Realistic electronic books

Veronica Liesaputra and Ian H. Witten

Department of Computer Science, University of Waikato
Hamilton, New Zealand

Abstract

We describe a software book model that emulates a range of properties associated with physical books—analog page turning, visual location cues, bookmarks and annotations—and, furthermore, incorporates many advantages of digital environments—hyperlinks, multimedia, full-text search, automatic identification of synonyms, cross-referencing of key terms with an online encyclopedia, and an automatically generated back-of-the-book index. Usability studies were conducted to compare performance using these books for various reading tasks with HTML, PDF and physical books. Participants completed the tasks more efficiently with the new interface without any loss in accuracy; they also preferred it.

Keywords: Within-document navigation, Within-document search, Usability studies, Wikipedia, Electronic book

1. Introduction

The electronic book industry is burgeoning. Libraries and other content providers are digitizing documents to allow readers around the world to access them online. The Open Content Alliance and Google are creating collections of tens of millions of volumes accessible through the Internet Archive and Google Books respectively. Manufacturers compete to produce reading devices that are light and mobile, with high resolution displays and many other features—Barnes and Noble's Nook, Amazon's Kindle, and the Sony Reader are typical examples. However, neither the underlying document representation nor interaction with electronic books has evolved significantly—apart from a few superficial visual effects. As a result, although users are increasingly willing to read documents sequentially on-screen, most still prefer to print them when they want to study them intensively (Nicholas et al., 2008; Gorman & Crawford, 1995; Crawford, 1998).

The use of books to store knowledge has a long tradition—books are arguably the most mature user interface ever devised for presenting information. People have acquired a variety of strategies and manipulation techniques, such as using fingers as bookmarks or riffling through pages while browsing (Severinson et al., 1996; Kerr, 1986). Compare a medieval book with a modern one: the same principles are still at work. Readers need only learn these conventions once, and can use them for the rest of their lives.

Current electronic document systems offer great added value over paper books: authors can revise information quickly and incorporate hyperlinks and multimedia, and readers can locate words or phrases through full-text search. However, electronic book applications fail to provide adequate cues about the reader's location in a document, and do not support interactions that are readily achieved with physical books without disrupting the ongoing reading activity (Dillon, 1992; O'Hara & Sellen, 1997; Conklin, 1987; Edwards & Hardman, 1989; McDonald & Stevenson, 1998). These interactions are exactly the features that physical books have evolved to support.

Opinion on the importance of emulating the appearance of printed books in the electronic environment is divided. Some researchers argue that it is unneces-

Reader's engagement / Number of documents	**Passive reading**	**Active reading**
Single text	Enjoying a novel, reading a poem	Studying a textbook, reviewing a proposal
Multiple text	Surfing the Web, reading e-mails	Researching a problem, surveying a topic

Table 1: Categories for reading activities (Adler & van Doren, 1972; Schilit et al., 1999)

sary to paginate documents or maintain their printed page layout (Rowland et al., 1995; Nielsen, 1998b; Kol & Schcolnik, 2000; Shneiderman, 1998; McKnight et al., 1991). They believe that with training and practice, users will no longer rely on the book model to locate information. Others contend that users are dissatisfied with digital document systems because they do not embed familiar book metaphors (Benest, 1990; Henke, 1998; Landoni et al., 2000; Crestani & Ntioudis, 2001; Barker et al., 1994). Their studies show that all participants, regardless of computing experience, find the book metaphor easy to use and understand, and that people can utilize their familiarity with paper documents to navigate and obtain information from digital text.

In order to investigate whether digital books can capitalize on reader experience with physical documents to enhance electronic reading, we have developed a software model called Realistic Books that is designed to combine the advantages of physical and electronic documents, and implemented a way of creating such books automatically from any HTML or PDF file. We have also compared performance using the new interface with existing electronic reading interfaces, and with the original physical form itself—the humble book.

Over nearly three decades, researchers have developed various electronic document representations, such as Book Emulator (Benest, 1990), XLibris (Schilit et al., 1998), and 3Book (Card et al., 2004b,a). Most are claimed to have glowing potential, but their proponents rarely make objective empirical comparisons of user behavior and performance with conventional representations. Because these new representations are not available for others to use, we cannot compare their performance with that of Realistic Books, nor with HTML, PDF and physical books.

After defining the scope of the investigation, this paper reviews the historical development of document systems and their effect on reading strategies and behavior, in order to identify, in Section 4, aspects of both printed and electronic documents that might enhance reading performance. Section 5 describes how documents are presented as Realistic Books, and the functions that are supported. Section 6 explains how Wikipedia can be used as an encyclopedic knowledge base to automatically identify key terms and phrases in book text, provide on-demand definitions, and search related pages in the book. Finally, we evaluate system responsiveness and present the results of user studies that compare people's browsing, searching and annotation performance and behavior in simulated books with the HTML, PDF, and physical representations.

2. Scope of the investigation

Despite its quotidian familiarity, reading includes a range of activities, is done for various purposes, and is embedded within many other document-based activities. People read for different reasons and in different ways: locating information, checking facts, acquiring new knowledge, analysing text, recreation. Studying a textbook is not like reading a science fiction novel.

Adler & van Doren (1972) and Schilit et al. (1999) characterise reading along two dimensions: the

reader's engagement with a text, and the number of documents involved, as shown in Table 1. Reading a single document involves within-document navigation, while reading multiple documents also involves filing, sorting and inter-document navigation. However, multiple-text reading activities lie outside the scope of this research.

Reader engagement with a text ranges from passive to active (Adler & van Doren, 1972). Passive reading is often associated with reading for leisure or entertainment, as when reading novels, poems and e-mails. It requires less thought and effort than active reading. Active reading involves not only reading the document, but also critical thinking, learning and decision making. For example, while a student is studying a textbook, he or she might mark some pages of the book, and add their own commentaries to the text.

Users have a different experience when reading online documents (Gould et al., 1987; Muter & Maurutto, 1991). Most physical books are portrait rather than landscape oriented. This is the opposite to that of monitors for PCs or laptops, and consequently readers cannot comfortably fit the full page on the screen (Dillon, 1992). Users must move the document around to be able to see all components of the page. If the document is scaled to fit the screen, its content may not be legible. Furthermore, the act of zooming in and out is commonly rated as cumbersome and difficult (Waycott & Kukulska-Hulme, 2003). One solution is to allow readers to alter the font size so that the text is large enough to be read comfortably for long periods of time. However, this changes the length of the document and users find it distracting to move away from their current text to obtain navigational information such as section headings (Dillon, 1992). Small screen displays lead to more navigation, and are not suitable for lengthy material (Marshall & Ruotolo, 2002; Waycott & Kukulska-Hulme, 2003).

Form factor is a particularly important issue for reading appliances. Size, type and quality of a computer screen affects reading speed, number of pauses, concentration time, and the decision to read or skip material (Burbules, 1998; Muter, 1996). Desktop and laptop computers provide users with larger workspaces, but suffer from issues such as heaviness, fatigue and eyestrain associated with computer screens. Current digital display media, such as Kindle, Sony Book Reader and iPad, are compact, lightweight and facilitate stylus or gesture based input. However, their small screens and lack of good keyboard support make them unsuitable for active reading tasks (Ryman, 2010; Morris et al., 2007; Marshall & Ruotolo, 2002). Furthermore, devices that utilize the E-Ink Vizplex electronic paper display require extra lighting for reading in darkness, offer only black-and-white displays, are not suitable for displaying rich multimedia and interactive materials, and readers cannot flip quickly back and forth between pages (Genuth, 2007; Martinez, 2010). Because of their limited capability, many people prefer to use a laptop or desktop computer (Williams, 2000).

There has been much research and commercial interest in developing the reading devices that are affordable, portable, durable, easy to use, readable, comfortable, have large storage facilities and a long battery life, can be connected to the Internet, and support good keyboard, stylus and gesture-based input. For example, Chen et al. (2008), Hinckley et al. (2009), Ajmera et al. (2009) and Ricker (2011) are creating dual screen displays; Universal Display Corp., Industrial Technology Research Institute, Polymer Vision and HP are experimenting with flexible displays that can be rolled up into the size of a small cellphone and rolled out into a big screen display when needed; Hinckley et al. (2010) evaluate a system that accepts multimodal pen and touch inputs; Immersion and Microsoft are working on shape-shifting display technology that allows sections of a touch screen under a person's fingertips to rise up or sink down; and eInk, Samsung and Qualcomm have developed colour electronic paper display. Unfortunately, the way in which electronic documents are presented in those new reading devices has not evolved.

There is enough evidence, both anecdotal and from formal user studies, to suggest that the usual HTML or PDF presentation of documents is not always convenient, or comfortable, for the reader (Kearsley, 1988; Parsons, 2001; O'Hara & Sellen, 1997;

Malama et al., 2005). Although a portable reading device with a touch-screen display can create a richer and more tactile experience, digital interactions with the documents are still interruptive and deflect users from their reading tasks (Marshall & Bly, 2005; Adler et al., 1998; O'Hara & Sellen, 1997). Furthermore, the lack of physical aids and standards in the text structures cause users to frequently become lost in the document (Conklin, 1987; Edwards & Hardman, 1989; McDonald & Stevenson, 1998). They do not know where they are or where to go next. Compounded by factors such as reading ability and prior knowledge, an electronic environment may increase the user's cognitive reading load and reduce his or her reading performance (Niederhauser et al., 2000; Gordon et al., 1988; Marchionini & Shneiderman, 1988; McKnight et al., 1989).

Given the remarkable historical success of the book form, and the fact that the superiority of HTML or PDF presentation is at least questionable, why are physical book models not more widely used in digital libaries? This paper asks whether this is purely technological or whether there are any proven usability issues.

Many people still undertake knowledge work on their standard desktop or laptop. Thus we investigate whether the book form provides the best document representation and interaction for a user to actively read a single document on a desktop or laptop. It is assumed that the subject's reading experience and performance will be further enhanced if the proposed book form were augmented with missing functionalities, such as the ability to share and collaboratively interact with multiple documents, or displayed on an ideal digital reading device. In the future, we will investigate whether the results of this study apply to touchscreen computers of varying size, like Apple's iPad, Acer's Iconia or Sony's S2.

3. Evolution of the book form

Are books adequately represented by their content, independent of their form? The book form conveys abundant information using well established conventions that everyone is accustomed to and can interpret intuitively. The conventions for book layout and design are the legacy of a long period of evolution. Their familiarity renders a book's graphical elements and functionality practically invisible. The "codex"—a book with pages that turn rather than a scroll that is read by unrolling—is one of the greatest inventions of all time, a technological innovation that significantly accelerated the transfer of knowledge. History shows that the form of the written medium dictated how the text in a document is organized and read.

3.1. Printed documents

The first material used for portable writing was the clay tablet of the Sumerians in 3500 B.C.E., shown in Figure 1(a). Because they were generally intended for administrative records, such as to help farmers remember the number of goats in a field, tablets were designed to fit comfortably in the reader's palm. Some documents comprised several tablets, perhaps kept in a leather pouch or wooden box in a predetermined order for sequential retrieval (Diringer, 1953).

Scroll. Around 2400 B.C.E., Egyptians began to write on scrolls made up of papyrus plants, stripped and connected with glue created from the sap of the plant (Baines, 1983). Figure 1(b) shows an example. At 135 feet, the Great Harris Papyrus, composed by King Ramesses IV, is the longest scroll ever found in Egypt (Dunn, 2006).

Initially few people could read, so public readings were common (Harris, 1989). Scribes wrote down text that was dictated to them and read it aloud on sacred occasions. Manuscripts were written in *scriptio continua* (Gaur, 1984), without capitalization, spaces, punctuation marks, or any other external indication of structure. Because scribes often knew the text they were transcribing by heart, they needed few visual aids: the script was only used as a cue sheet.

The use of scrolls was well suited to the continuous style of writing because, during oral reading, documents were always read sequentially, one frame at a time; the reader rarely jumped to different sections of the text. Nevertheless, reading a scroll is not a simple matter, for both hands are required to roll and unroll

(a) (b) (c) (d)

Figure 1: Ancient devices for portable writing: (a) Sumerian clay tablet (Choi, 2010), (b) scroll (Easley, 2008), (c) concertina (Fuller, 2006) and (d) codex (Drucker, 2006)

continuously. Similarly, to locate a specific passage it is necessary to spool back and forth. Frequent unrolling and re-rolling make scrolls liable to damage.

Concertina. By 400 B.C.E., reading was no longer a simple memory aid but a medium for interpreting and creating information. Although readership was still small and socially restricted, reading circles began to appear, and authors would exchange documents (Lichtheim, 1973). To help those with poor reading skills, scripts were written in *per cola et commata* style, the text being divided into lines that have coherent meaning and represent a phrase, clause or sentence (Wingo, 1972). This made it easier to search for a passage. Each line was treated as a small paragraph by projecting the first letter into the margin, and readers were expected to take a breath at that point. Punctuation marks and vowels were introduced to help people assimilate the text.

To overcome the limitations of scrolls, concertina or accordion books were developed. Intermediate between scrolls and codices, several sheets were glued into one, which was folded onto itself as illustrated in Figure 1(c), often with hard covers to protect the pages. People usually folded each segment of their concertina books at line breaks. This format was preferred over the scroll. When folded it resembles an ordinary book; readers can randomly access any page, without unrolling and rolling, to search for a passage, and read the document sequentially by flipping the folds. Unfolded concertina books are essentially scrolls, so those who are comfortable with scrolls can use them in the same way. When folded they are smaller than scrolls: easier to handle while reading and easier to store afterwards. Julius Caesar folded rolls into concertina books to send orders to his troops (Suetonius, 2000).

Codex. Later, the codex or book form, illustrated in Figure 1(d), was developed. Around 200 B.C.E., Greeks and Romans began to use wax tablets backed with wood, metal or ivory to write orders or notes—sometimes connected with cords (often looking like a three ring binder) (Diringer, 1953). Around 100 C.E. Romans substituted parchment, which was lighter, more durable, and resistant to insects and humidity, for the panels in codices.

By the end of the 4th Century C.E., the standard form of documents had largely changed from scrolls to codices, and concertinas were relegated to maps and posters. The codex became the preferred format because people could write on both sides of the material, and it provided easy access to any part of the text. Fixed layouts promote spatial memory and help readers locate information they have seen before.

Codex books were cheaper and more convenient than scrolls and concertinas. They used less material, were easier to read and store, and were portable and searchable. The format prompted innovations in information handling: text could be organized in books or chapters according to its content, transcending the limited capacity of a scroll (Reynolds et al., 1983). The only real disadvantage was that readers could not view more than two pages at once.

As the book's physical form and text arrangement became simple, clear, and more appealing to readers, around 900 C.E. reading started to change from

a public act to a silent, private one (Cipolla, 1969). With the invention of the moveable-type printing press, more affordable books meant that people no longer treated them as sacred objects. Passive, ritualized reading slowly evolved into a more active and individual form of reading. People read information in confidence and at leisure. They could cross-reference and compare with other books, and critically evaluate the information they read. They annotated and bookmarked text to help them understand, summarize and organize the material for later review.

3.2. Electronic documents

Curiously, the evolution of computer output has paralleled the development of the book format. Paper rolls were used for early printers—scrolls. Later, line printers used fan-fold paper, perforated so that it could be folded, boxed, and handled more easily—concertinas. Today, people print on pages and staple them together into books.

The parallel is not confined to print technology. Over the last decade many electronic document representations have emerged, and the evolution from scroll to codex can be discerned here too.

Scroll. Early text display monitors scrolled; so do webpages, originating in the early 1990s. Examples are legion, including text repositories maintained by Project Gutenberg and articles in Wikipedia. In this format, documents are represented as a long continuous page with horizontal and vertical scroll bars. Readers can usually change the page size dynamically according to the screen space available.

Concertina. Much document-oriented software, like Adobe Reader and Microsoft Word, represents documents as scrolled sequences of pages, like concertinas. Text is segmented into pages that preserve the logical structure, formatting and layout of the book. Readers can either scroll or page through the document, enjoying both the advantages and disadvantages of scrolling and paging.

Although LiquidText (Tashman & Edwards, 2011) and source code editors like Dreamweaver and Eclipse display a document as a long continuous page, they allow users to hide or collapse parts of the document to bring text from disparate areas of a document together or to hide all immaterial text. This is like viewing a concertina electronic book format without the ability to page through the document.

Stack of pages. Barnes and Noble Nook, Amazon Kindle, Sony Reader, XLibris (Schilit et al., 1998), Microsoft Reader, and Hyper-Text Book (Crestani & Ntioudis, 2001) all present documents one page at a time. This is like viewing a concertina electronic book format without the ability to scroll from one page to the next.

Codex. Many designers have begun to embed the book metaphor into the electronic environment. Documents are presented as double-page spreads with a stack of page edges on either side of the opened page. The appearance and handling of physical books can be simulated either in a 2D environment such as the Book Emulator (Benest, 1990), Visual Book (Landoni et al., 2000), BookReader (Internet Archive, 2007), Turning the Pages (British Library, 2006) and iBook (Apple, 2010), or in a 3D environment such as the WebBook (Card et al., 1996), 3D Book Visualizer (Chu et al., 2004, 2003) and 3Book (Card et al., 2004a,b).

4. Design considerations

Navigation and personalization in digital reading both present many usability issues: accessing and opening the book; knowing where one is in it; providing and accessing supplementary information; navigating within it; locating information by full-text searching or an index; and personalization with annotation or bookmarks. This section briefly reviews research findings on these—findings that are often contradictory.

4.1. Accessibility

Before beginning to read an electronic document, a viewer must be launched and the text retrieved. If network delays are involved, there may be a distractingly long lag before reading can begin (McKnight

& Dearnley, 2003). Even without a network the start-up delay may be significant: Nicholas et al. (2008) found that their participants dislike Adobe Reader for this reason.

Not every user has their own computer. Many share with others, and may not be able to install applications or even store files on the communal machine. Electronic book programs must compete with other applications for disk space, installation and upgrading effort. In many shared systems, only the most common applications are installed. Even with their own computer, users often shy away from downloading and installing new software.

4.2. Orientation

When readers of a physical book are interrupted, they simply resume reading at the point where they left off. If they happen to lose their place, the book provides a strong sense of context: much of the reader's field of view is devoted to orientation cues such as running headers, page numbers, section titles, and—importantly—the stack of page edges on either side of the reading position. Readers are subliminally aware of their position in the document and its overall length (Crestani & Ntioudis, 2001). Although some electronic formats display running headers and page numbers, and all show section titles, the stack of page edges is invariably absent—its function being performed by a scroll bar or page number.

With physical books, readers have a sense of how old a document is and which pages are read often. Visual features such as crease marks, decolorization and coffee stains are evident. With many electronic book applications, users rely on the document's metadata to determine when the document was created and last modified. Hill et al. (1992) tried displaying coloured marks in the scrollbar to simulate "read wear" by showing how often sections of the document have been read.

Although no reading applications have utilized these markings to show history of use, the idea has inspired several researchers to use markings as navigation cues to help users quickly revisit sections of the document (Laakso et al., 2000; Bjork, 2001). When a reader or an application bookmarks a certain part of the document, a coloured mark is placed on the scrollbar at a position that corresponds to the bookmarked text's position in the document. Clicking the mark takes the reader to that point. Several software development environments use this idea for marking code errors or comments, and Web browsers use it to mark the location of matched search terms.

Readers of electronic documents may prefer scrolling on the basis of familiarity, because it is perhaps more common and better established than paging. Mills & Weldon (1986) found no significant difference between the two. On the other hand, Parsons (2001) found that users were dissatisfied with the scroll format. Scrolling weakens spatial relationships and makes it hard to read documents non-linearly. Also, for long documents, moving the scroll-bar slider up or down just one pixel can completely change the screen contents.

For paginated documents, systems generally present readers with the current page number and total page count. Not all users benefit from this information, or even notice it. Marshall & Ruotolo's (2002) subjects preferred analog visual feedback, like page edges, to inform them where they were and how long the book was. Digital numerals were less helpful because users could not easily visualize this information.

4.3. Text comprehension

The more readers know about a text, the more likely they are to understand it. Their knowledge might be about the language, the world in general, or the subject domain. Comprehension can be increased by providing background information that explains relationships between ideas in the text that are left unstated in the main body of the document (McNamara et al., 1996). If such information appears in traditional books, it is generally found at the back: appendices, references, glossaries and indexes.

With electronic documents, comprehension can be enhanced by hyperlinking words and phrases to a glossary (Landoni et al., 2001). Readers who are unsure of the meaning of a term can click it, and

straightaway see its definition. Similarly, people often prefer to read foreign texts online rather than on paper because they can immediately consult dictionaries or language translations (Crane, 2002). Links can also be used to call attention to important items of text. For example, many online newspapers link key concepts to pages that explain them.

On the other hand, the use of active hyperlinks to supplementary material can be confusing and distracting (Wilson et al., 2003; Marshall et al., 1999). Readers might find it unclear which hyperlinked items lead to glossary entries, references, other sections of the book, or external documents. Those who already understand all the terms in the text need not consult glossaries, and might find superfluous links distracting.

4.4. Within-document navigation

Few electronic book applications implement realistic page turning. In order to proceed to the next (or previous) page, users click buttons that step through the document. Page turning is invariably a discrete act rather than an analog one (Benest, 1990), taking the form of blanking the screen and drawing the new page on it. This may be efficient, but readers briefly lose contact with the text, which interrupts their interaction.

During Marshall & Bly (2005)'s navigation study, participants complained about their inability to turn pages in the electronic book. Without this, they could not continue to see the text on the current page while peeking ahead to the next for a foretaste of what was in store. The new version of Adobe PDF and iBooks incorporate an interactive page-turning mechanism; iBooks even include non-linear surface distortion so that letters appear to curve around the page as it is turned.

However, according to McCusker (1998), analog page turning is ineffective and inefficient. The feeling of physically turning a page can be preserved simply by animating the page-turn motion when users click the *previous* or *next* button. This animation is sufficient to maintain reading continuity, psychological contact with the text, and inform readers of their direction of travel. Furthermore, study of people's behavior while reading with devices like Kindle or Sony Reader show that participants like the fact that they do not have to turn pages (Schcolnik, 2001).

4.5. Searching

Physical books provide a table of contents and back-of-the-book index to be used as quick points of entry to the document content, and as a map for navigating through the book. Electronic documents usually have a table of contents and a full text search function: they rarely include back-of-the-book indexes. Instead, users type a word or phrase to obtain the desired information by searching. However, people use a variety of words to refer to the same thing.

This is exactly the problem faced by subject indexers when deciding on terms for a back-of-the-book index. Ideally, indexes should contain all terms that readers might conceivably use when seeking information in the book. Nevertheless, entries are not chosen based on the infinitude of possible search terms. Good indexes bridge the author's perspective of the subject and the terms describing the book's contents with the range of keywords that most readers might use in an enquiry.

To further help users locate relevant information, given a book for which a back-of-the-book index available, 3Book reorganised the index every time people enter a query (Card et al., 2004a). When users entered a search term, they see only the index entries that are relevant to it; all other entries are hidden. Their user study showed that participants were faster in finishing their tasks and more accurate in their answers using the reorganised index rather than its paper version, regardless of their familiarity with the subject matter. Although users preferred the 3Book's back-of-the-book index over the paper version, it was not clear whether they would still use it if a full-text search function were available.

Ryan & Henselmeier (2000) compared user performance when finding information in electronic books with and without back-of-the-book indexes, and found that indexes helped participants locate information more quickly and accurately. Nevertheless, they rated full text search more important than

the index—in fact, few subjects mentioned the index as a desirable electronic book feature. They felt that the search facility was slightly more engaging than the index.

4.6. Personalization

Online annotation and bookmarking systems are meant to give electronic documents the same personalization facilities as paper ones. Some, such as ScreenCrayons (Olsen Jr. et al., 2004) and Adobe PDF, allow readers to write wherever they want without selecting text first, while others, such as Annotea (Kahan, 2001) and Kindle, only permit annotations once text is highlighted. Some display annotations and bookmarks on the same page; others on a separate page. Some support free-form drawing, highlighting and text comments.

Yet in contrast to the richness of personalization practices in printed books, readers rarely personalize online documents. Instead they prefer to print, and annotate the paper version (Liu, 2005). Online annotations are distracting and require more effort and practice than a pencil or highlighter (O'Hara & Sellen, 1997). They are not seamlessly integrated with reading: users first select a button to start annotating and then draw with a mouse or type on the keyboard. However, hand-held reading appliances with touch screen displays free up one hand, and can imitate the ease and flexibility of annotation on paper (Marshall et al., 1999).

People prefer to type their annotations than using a stylus to write them (Morris et al., 2007). Readers tend to write in a bigger font on an electronic reading surface than when they are writing on a paper, i.e. they need a bigger space for their annotation. They also prefer to type because this makes their annotations searchable, editable and part of their contents could be copied to other applications.

Buchanan & Pearson (2008) report that readers prefer bookmarks to appear as tabs beside page edges, instead of being listed in a drop-down menu as in web browsers or displayed in a panel alongside the page as in Adobe's PDF Reader. Tabs provide stronger spatial cues about the bookmark's position relative to the reader's current position in the document. Nevertheless, participants also expressed concern about the limited screen space when bookmarks were displayed as tabs or a sidebar. Listing bookmarks in a menu or sidebar only provides readers with the associated page location and title, which is insufficient to facilitate many forms of information access (Cockburn & Greenberg, 1999).

4.7. Three-dimensional environment

WebBook (Card et al., 1996) and 3Book (Card et al., 2004a,b) represent electronic documents using a three-dimensional model of a physical book. Although these authors describe the possible reading activities that their system can support, they did not test its effectiveness in supporting those activities. Therefore, it is not clear whether any of their features, which include a conceptually reorganized back-of-the-book index, free-form annotation and the use of a book metaphor in a three-dimensional environment, will significantly increase reading performance.

Expert views on the utility of a three-dimensional over a two-dimensional environment as a visualisation technique are divided. Spatial organisation of information supports rapid retrieval by users, and some believe that three dimensions can be used to exploit human spatial capabilities (Czerwinski et al., 1999; Tavanti & Lind, 2001). Others believe that even though spatial organisation can improve the user's ability to retrieve information, there is no significant difference between the usefulness of the two environments (Chen, 2000; Swan & Allan, 1998; Sebrechts et al., 1999). These people also believe that a user's experience and topic choice influence their recall and precision more than their spatial ability.

Studies were carried out by Cockburn & McKenzie (2001, 2002) and Cockburn (2004) to evaluate users' spatial memory effectiveness in both environments. They found that most users thought the three-dimensional interface was more cluttered and less efficient. The additional flexibility or freedom to move around did not help them in retrieving information. They became confused, and their ability to retrieve information deteriorated. They did not navigate well

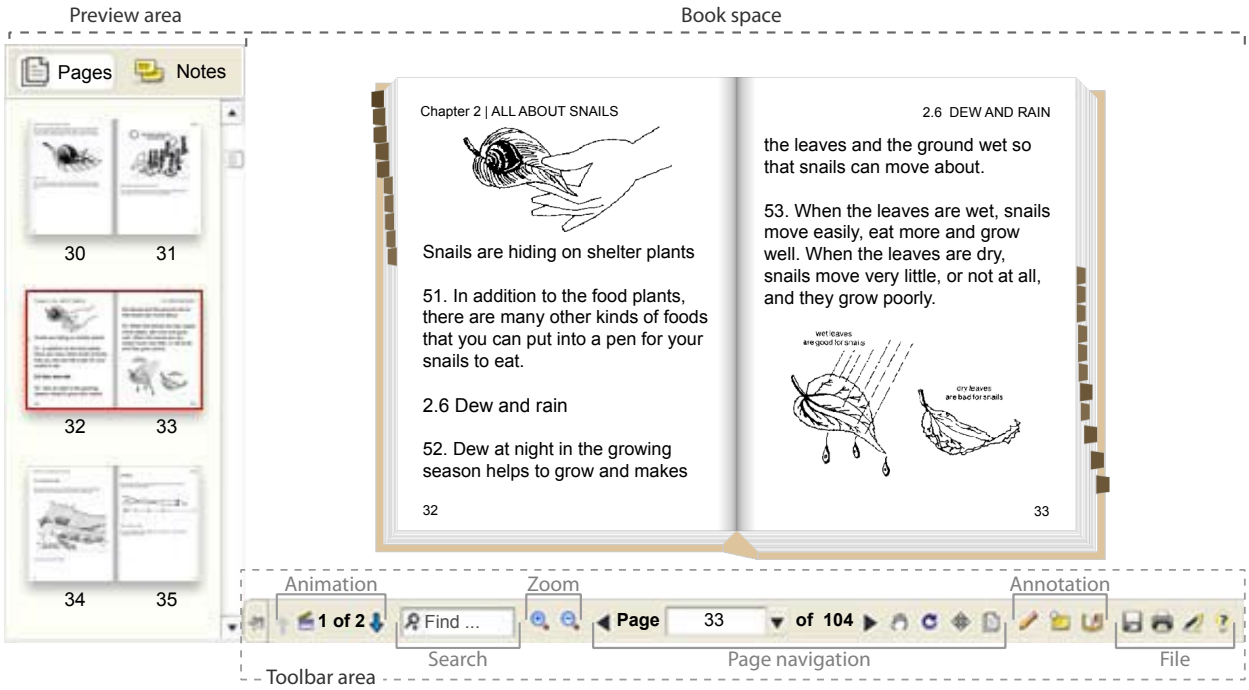

Figure 2: Reading a Realistic Book (readers can suppress the widgets surrounding the book space)

in this environment unless they could benefit from previous experience with a three-dimensional system. The best environment for a particular system depends on the application and the interface.

Nielsen (1998a) states that the main problem with three-dimensional visualisations is that current interaction techniques were designed for two-dimensional manipulations. With three dimensions, users have extra controls that allow them to manipulate their view. They are required to learn how to navigate around the environment and how to manipulate the primary model. Navigation gets in the way of their primary task. Some guidelines to alleviate these problems are to simplify and minimise navigation steps necessary to complete a task, and to avoid any unnecessary visual clutter.

According to Dillon (2000), these sceptical views stem from users' unfamiliarity with the interfaces. Once readers gain experience, they start to notice the advantages of the environment. Three-dimensional metaphors can produce an intuitive representation of document metadata, and allow a straightforward analysis of the document collection. Information space visualisation techniques provide an ideal setting for interactive browsing (Dumas et al., 2002). Three-dimensional environments have practically taken over video gaming, and people usually find it easy to navigate without any training. It seems that in the future users will get used to the idea of navigating simulated space.

5. Realistic Books

To investigate whether book models with realistic page turning offer measurable advantages over physical books and other electronic forms, a lightweight Adobe Flash-based application, called Realistic Books, was constructed. Flash was chosen because it is widely installed, and minimizes both download time and compatibility issues (Millburn Survey, 2010). Realistic Books communicate with a PHP application server to save the reader's annotations.

A key challenge when simulating physical books is modeling the simple act of turning a page. Liesaputra & Witten (2009) review several methods of varying complexity. A 2D technique can be utilized to create a sufficiently natural simulation of an analog page turn that is responsive, scalable and can be easily implemented: this is what the Realistic Books use. However, 3D models are required to simulate

turning a block of pages together, or riffling quickly through many pages.

An informal usability study with the 3D Book Visualizer (Chu et al., 2004, 2003) showed that participants required a significant amount of learning to know how to use the extra three-dimensional controls for panning, zooming and rotating, to manipulate their document view. Even experienced computer users found these controls confusing and error-prone. Because page turning simulation with 3D models is far less responsive than with the 2D technique, Realistic Book system uses a 2D page turning model.

Once the page-turning act has been simulated, many more details must be attended to before arriving at a compelling model of a realistic book (Chu et al., 2004). These include modeling the interaction between the pages and the covers, the effects of metadata such as bookmarks, and tools for supporting various associated activities such as searching and annotating. These details are crucial to the realization of a comprehensive electronic book application.

Figure 2 shows the Realistic Books interface. The screen is divided into three areas: the book space, the reader's tools and the preview area. Most of the screen is dedicated to the book itself—many readers will choose to suppress the other elements, and book designers can arrange for their books to open this way. By default, only the book space and the reader's tools are visible. This section describes how electronic documents are presented as books and outlines the reading services that are provided to satisfy the generic requirements of readers. These services are designed to follow the book metaphor and the features of normal book consultation as closely as possible.

5.1. Geometric book model

A book can be presented closed, with the front or back cover in view, or open, with a double-page spread in view. The implementation uses parallelograms to represent the topmost pages, the stack of page edges, the front and back covers, and the book's spine. Although curved shapes may appear more realistic, flat parallelograms enhance readability and ensure rapid page-turning (Liesaputra & Witten, 2009).

When the book is open and a page is being turned, users see more than the topmost double-page spread: part of the back side of the page being turned and part of the page underneath are also visible. However, in all the book's possible poses, parts of no more than the two topmost leaves on either side of the double-page spread are ever visible to the reader.[1] For all other leaves, readers can only see the edges. This makes it unnecessary to store the entire book's content in main memory: only the recto and verso page contents of the two topmost leaves are needed. All pages are assumed to have the same background color, so horizontal and vertical page-edge textures can be used to represent the left and right page blocks.

For example, when the book is closed only the front cover is shown, but the front endpaper (verso side of the front cover) and page 1 are also loaded. When the reader turns the front cover, as illustrated in Figure 3(a), these are all that is needed; when the turn is complete, the contents of pages 2 (verso side of page 1) and 3 are loaded. Similarly, when the reader turns the left-hand leaf of an open book to the right-hand side, as shown in Figure 3(b) where page 6 is being turned back to reveal pages 4 and 5, the recto and verso content of the leaf being turned (pages 5 and 6), the recto content of the leaf below the turned page (page 4), and the verso content of the right-hand page (page 7), are all that is needed to simulate the page-turn. Once it is complete, the contents of the appropriate pages (in this case, pages 2 and 3) are loaded in preparation for a possible subsequent turn, and the position of the left and right page blocks are updated.

Each element of a Realistic Book, such as page size, turning speed, and text margin size, is controlled by markup specified in a template file for the book, which is expressed in XHTML. By default, the page

[1] Where necessary for clarity, we use "leaf" to refer to a piece of paper, which has front (recto) and back (verso) sides. This contrasts with the term "page," which denotes a single side of the paper.

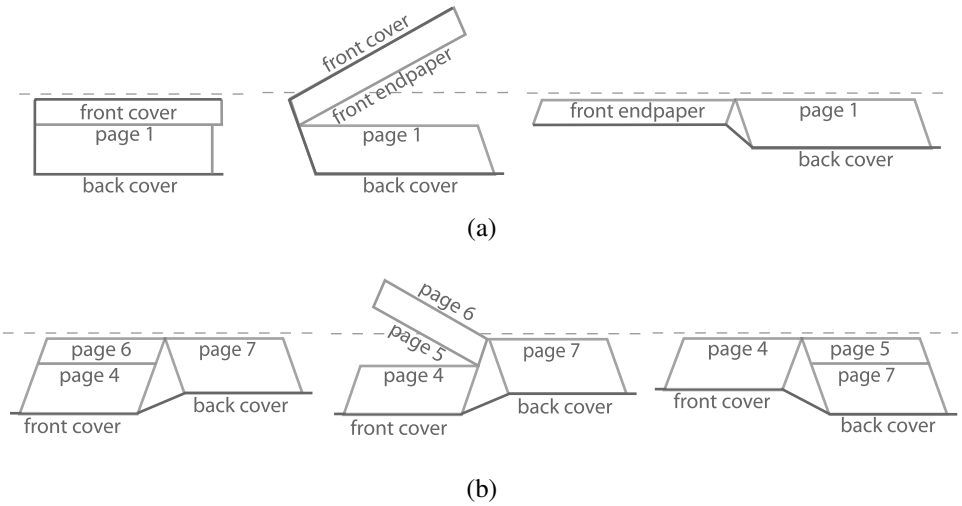

Figure 3: The book model: side view when (a) the front cover and (b) a left-hand page is turned.

aspect ratio is set to that of an A4 page. For presentations, any given page can have several versions or "overlays," and the arrow keys move between overlays.

The thickness of the left-hand block of pages is the number of leaves on the left side of the book, and similarly for the right-hand block. The width and height of each block's base is the width and height of an individual leaf. Two parameters specify the vertical and horizontal offsets between each leaf in the block.

For a hardcover book, the cover protrudes beyond the other pages by a parameter called *BookMargin*; thus the cover's width and height are given by $PageWidth + 2 \times BookMargin$ and $PageHeight + 2 \times BookMargin$ respectively. The book's spine is drawn by connecting the front and back covers, and its color is set to the average color found around the border of the front cover.

By default, the table of contents, list of figures, beginning of each chapter and subsection, and back-of-the-book index, are all bookmarked. Bookmarks are tabs placed in predefined positions along the right or left page edges. The first is placed at the top, the second slightly below it, and so on down the page. The height of each tab is determined by a parameter that represents the greatest number of tabs that can fit along the edge of the page without overlap. Once this number is reached, the starting position is reset to the top of the page.

Although all bookmark tabs have the same shape and height, their width and color varies depending on type: top-level sections have darker and wider tabs than lower-level ones. To ensure that each tab matches the color scheme used in the book, by default the base color of bookmarks is calculated based on the book's cover and page color.

5.2. *Orientation and text comprehension*

A notable problem for online documents is homogeneity: their text tends to all look the same, making it difficult for users to get a sense of location (Nielsen & Molich, 1990). Realistic Books provide context and orientation clues through the table of contents, typographic cues, running heads, page numbers, page edges and bookmark tabs.

Each element of a Realistic Book's front matter begins on a right-hand page. The table of contents and list of figures can be automatically generated from the book's metadata. Each entry is hyperlinked: clicking it opens the book at the relevant section or figure. Page numbers indicate the location of entries and the length of sections. Pages in the front matter are numbered by roman numerals (except for blank pages).

Next comes the book's main body. Pages can contain text, hyperlinks, images, video and audio. Fig-

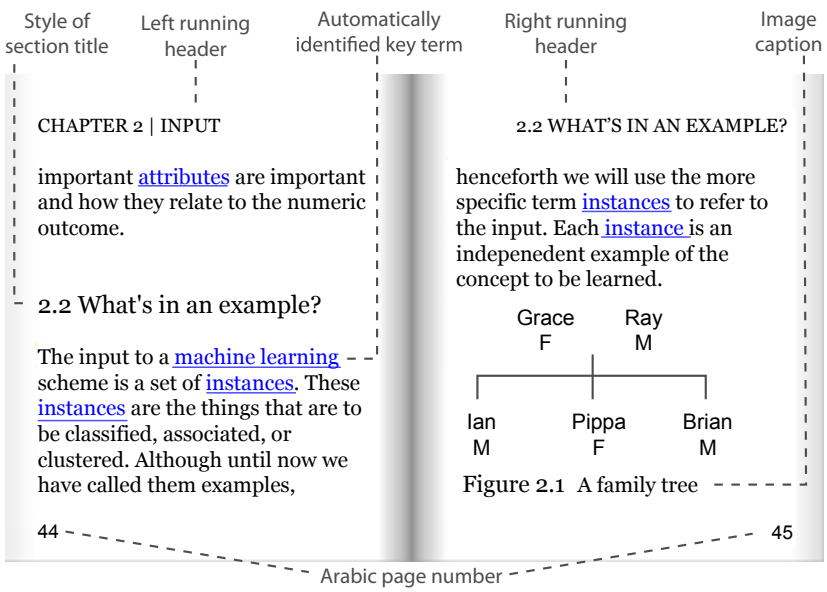

Figure 4: Typical page layout in a Realistic Book

ure 4 shows a page created using the default settings. Each chapter starts on a right-hand page. Headings and subheadings are rendered in boldface, in a larger font than the main text. Chapter titles are centered on the page, and section titles are left-aligned in a smaller font. Image captions are shown beneath the image in smaller letters. Running headers on each left-hand and right-hand page show the current chapter and section titles respectively. Running headers are omitted from the first page of each chapter. Starting from the main body, pages are numbered by arabic numerals (except for blank pages).

The book's content is processed to identify key phrases and their synonyms and hyperlink them to corresponding articles in an online encyclopedia. This process is explained in Section 6. Hovering over a term reveals its definition; clicking it opens its encyclopedia article in a web browser. Back-of-the-book indexes may be automatically generated from these terms.

The stack of page edges on both sides of an open book indicate visually the length of the document and the reader's position in it. Users can easily see where they are, and acquire incidental knowledge of the location of information by relating it to its physical position in the document. Similarly, bookmarks are shown as tabs that give visual cues to the structure of the book—an ever-present table of contents.

5.3. Within-document navigation

The pages in Realistic Books may be flexible or rigid, as illustrated in Figure 5. Flexible pages behave like ordinary paper; rigid ones behave as though they were made from stiff cardboard. The 2D peeling and shearing techniques described in Liesaputra & Witten (2009) are employed to simulate the act of turning flexible and rigid pages respectively. Simulating physical page turning allows readers to retain contact with the text, because the interaction is continuous; also, they can peek ahead to get a foretaste of what is in store.

Readers grasp a page by pointing at a corner—or anywhere along the top, right, or bottom edge—and depressing the mouse button. The page follows the motion of the cursor. It can be moved anywhere, within the physical constraints imposed by not tearing the paper, and the visual details follow instantly. When the button is released, the page either falls back to its original position or continues to turn until it lies at rest in the turned position, depending on whether or not the page has passed the center of the book. Although better control can be obtained by using a touch-panel, even with a mouse all this feels perfectly natural, as confirmed by the user evaluation reported in Section 8.

The system provides alternatives to the mouse ges-

(a)

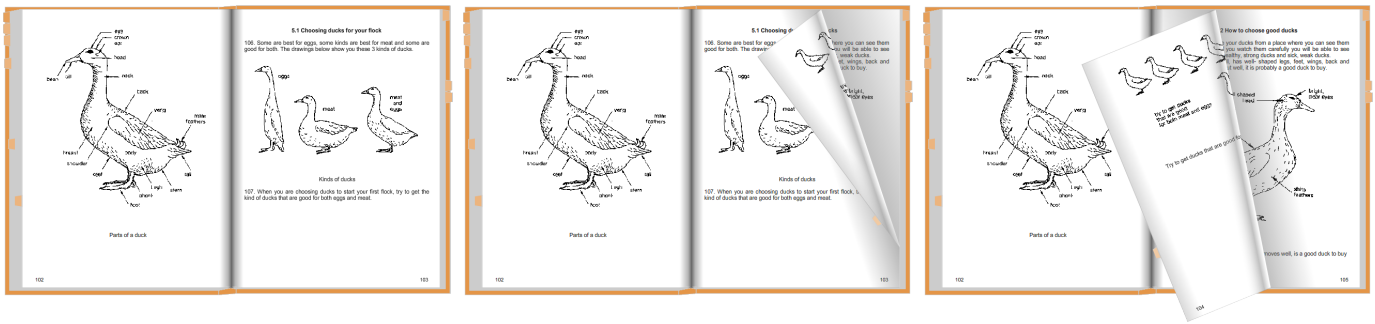

(b)

Figure 5: *Raising Ducks: How to Begin*: (a) opening the front cover, and (b) turning a page

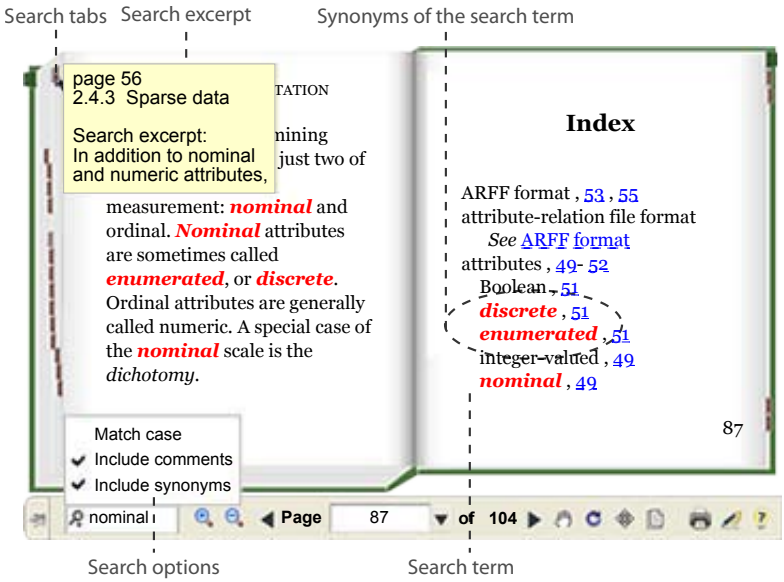

Figure 6: Text search interface

ture for turning pages. A right-hand page turns when it is clicked, the right arrow key is pressed, or the *next page* button on the toolbar is clicked; analogous actions work for left-hand pages. A *keep-flipping* button turns pages continually; this is stopped by clicking the book or pressing the escape key. An *options* button allows readers to select a suitable flipping speed.

As with physical books, readers can pick any page

edge and open the book there. Hovering over the page edges reveals the page number as mouseover text. When the book is open, the inside border of both covers is visible; clicking them closes the book (in either direction). Users can jump to a specific page by selecting or typing the page number in the toolbar (recall from Section 5.2 that pages in the front matter are numbered by roman numerals).

The *options* button allows users to display bookmarks as tabs that protrude from the pages to mark the beginning of each chapter and section, or, alternatively, the location of figures or annotations. This yields an ever-present table of contents, figures or annotations. Hovering over a bookmark reveals its page number and the chapter/section title, figure caption, or an excerpt from the annotation. Clicking it opens the book at that position.

It is often useful to be able to return to previously visited locations. Whenever a user performs a non-sequential navigation, the previously visited page is bookmarked. Clicking it turns the pages back to that location.

5.4. Text search

A prime advantage of electronic documents is the ability to search their contents. The Realistic Books search mechanism is modeled on Emac's incremental text search (Stallman, 2000). As users type terms into a search box in the toolbar, the program performs partial matching on the entire book. The results are narrowed down as each subsequent key is pressed.

When matches occur, all bookmark tabs except for those pointing to the table of contents, list of figures, back-of-the-book index and the previously visited page are hidden and replaced by tabs on all pages containing the search terms. Mousing over a tab shows the page number and an excerpt from the first matching fragment on the page, as illustrated in Figure 6. Clicking the tab opens the book to that page, with matching terms highlighted. Pressing the enter key takes readers to the matching page closest to their current position.

There are three search options that can be set using the menu shown at the lower left of Figure 6: *match case*, *include annotations* and *include synonyms*. When keyphrases have been identified using the method described in Section 6, the search process consults a table to find all synonyms of the query that occur in the book, and includes these in the search. These synonyms can be words or phrases. For example, a search for "Obama" will automatically find pages containing the phrase "44th President of America"—and vice versa.

5.5. Personalization

Realistic Books support two kinds of personalization: bookmarks and annotation. According to O'Hara & Sellen (1997), the act of annotating should not disrupt the reading activity, and readers should be able to annotate directly on the page in a way that stands out visually from the printed text.

In Realistic Books, all personalizations are superimposed upon the page without altering the original copy of the book. Bookmarks and annotations can be created at will and removed when no longer required. Readers do not need to press any buttons to switch from one to the other. Changes are automatically saved every five minutes and when the reader exits.

Clicking the *add bookmark* button automatically inserts a new bookmark on the right-hand page and places it in the first available position along the page edge. People can type comments and scribble; they can also insert images, resize, move or delete their annotations. They can specify the formats of the line and the text, and switch between a note's compressed and expanded form.

Comments can be inserted by moving the mouse pointer to where the note should appear and simply starting to type. Alternatively, selecting a portion of a page's text automatically highlights it, and comments can be inserted by clicking the highlighted text. To create a free-form drawing, users press the *pencil* button and draw anywhere on the page by depressing the mouse. Clicking the *hand* button or the escape key leaves the drawing mode and returns to the reading mode.

6. Enriched document representation

The information explosion powered by the World-Wide Web has led to a huge increase in skim-reading (Goldsborough, 2000). People scan documents for keywords, tables and illustrations in order to obtain a quick overview of their contents, and only read in depth when something catches their attention.

The background material that a document should provide obviously depends on the reader's prior knowledge (Shneiderman, 1989). To allow people with different backgrounds to understand a book, authors often provide glossaries and back-of-the-book indexes. Glossaries explain important terms that readers may find unfamiliar. Indexes show where each topic is mentioned, how it relates to the other topics, and what other terms could be used to describe it.

Studies by Furnas et al. (1987), Gomez et al. (1990), Good et al. (1984) and Buchanan & Pearson (2008) indicate that back-of-the-book indexes are redundant if the text is electronically searchable, particularly when user queries are automatically expanded to encompass synonyms and morphological variants of the search terms, and spelling errors are corrected. Well-established search engines already use these techniques to increase search quality.

Realistic Books provide for the automatic addition of a rich suite of reader aids and links to background material:

1. navigational hyperlinks
2. definition popups
3. synonyms
4. back-of-the-book index.

Each entry in the table of contents and list of figures is a hyperlink that points to the beginning of a section or a figure, generated from the section and image tags in the book's template file. These are the navigational hyperlinks.

The remaining three reader aids are generated by utilizing the online encyclopedia Wikipedia as a source of definitions, key terms and synonyms. Definition popups function like a glossary. The system automatically identifies certain key terms in the running text and displays a short description when readers mouse over them. Clicking on a key term takes the reader to the Wikipedia article that explains it. All key terms mentioned in the book are collected and listed in a back-of-the-book index. The index also includes synonyms of these terms; the synonyms are used in the electronic search function.

Key terms correspond to Wikipedia article names, and synonyms for a key term are found automatically by examining Wikipedia "redirects" or hyperlink anchor text to that article. For example, in Wikipedia the terms *Obama* and *Pres. Obama* redirect readers to the article entitled *Barack Obama*, and Realistic Books therefore treat these three terms as synonymous. When a reader types one of them in the search box, all locations in the book of that term and its synonyms are returned. Moreover, the back-of-the-book index will include entries

Barack Obama
⋮
Obama, Pres. Obama, *see* Barack Obama

(assuming that no other terms intervene alphabetically between *Obama* and *Pres. Obama*). This mechanism also handles spelling errors and common morphological variants.

We first provide a brief overview of Wikipedia and how its contents are organized, and then show how it can be used as a resource to automatically generate definition links, synonymous terms, and a back-of-the-book index for any textual document.

6.1. Wikipedia

Every article in Wikipedia defines and explains an entity, event or concept. The text of articles is studded with manually defined links that connect salient terms to other Wikipedia articles that explain them, offering users a deeper understanding of the topic. Articles are assigned to at least one category, and categories can be assigned to other categories, creating a graph organized into a loose hierarchy.

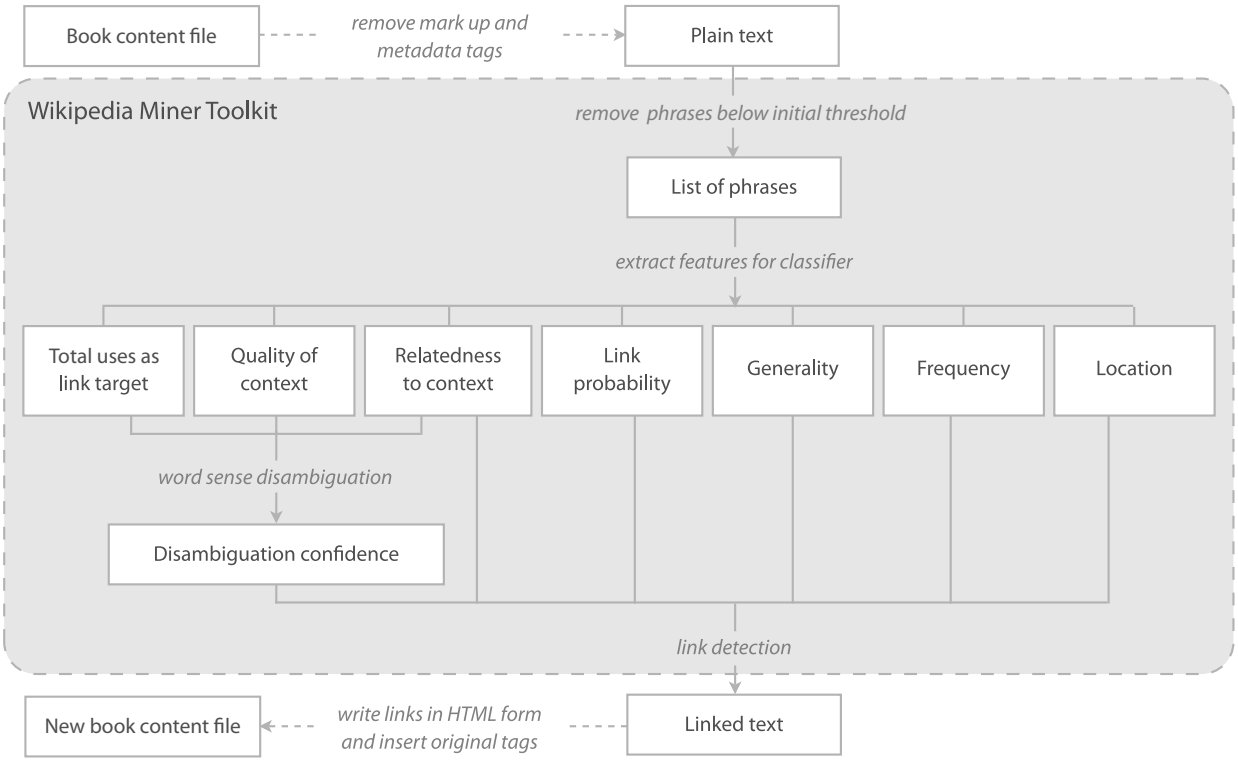

Figure 7: Automatically linking a book to Wikipedia

In natural language, the same concept is often referred to by different names—for instance, *magazine* is also referred to as *quarterly*, *serial magazine*, *scomparto*, amongst others. The most commonly used term (*magazine* in this case) is chosen as the Wikipedia article title, and readers are automatically redirected to it when they use one of the alternative names. The same technique is used to handle capitalization and spelling variations, abbreviations, synonyms, plurals, technical terms, and common variants. Our software also uses the anchor text of links in Wikipedia that point to other articles as an alternative way of referring to the target article.

Conversely, in natural language, any given word frequently has more than one possible denotation. For example, the word *Magazine* can refer to a periodical publication, a place to store ammunition, a light-tight chamber for film in a camera, an English rock group, or an Argentine TV channel. When the author of a Wikipedia article mentions a key term, he or she adds a hyperlink that points to the appropriate Wikipedia article—which depends on the context in which the term occurs. When the Realistic Books system identifies key terms in text, automatic disambiguation is essential.

6.2. *Linking documents to Wikipedia*

The process of extracting significant terms from text and augmenting them with links to appropriate Wikipedia articles is called "wikification" (Mihalcea & Csmoai, 2007; Zhou et al., 2009). Realistic Books utilize the Wikipedia Miner toolkit developed by Milne & Witten (2010) to automatically create definition links. This toolkit performs disambiguation using machine learning techniques that use the manually disambiguated links in Wikipedia as training data.

Figure 7 illustrates the process of generating definition links; full details, which are rather complex, can be found in Milne & Witten (2010). HTML and metadata tags are removed from the book's content and the clean text is passed to Wikipedia Miner. The wikification process begins by gathering overlapping word n-grams from the text and comparing them with a pre-computed label vocabulary that contains all Wikipedia article names, redirects, and hyperlink

anchor text fragments; it then disambiguates the resulting terms. For each possible sense of a term, the probability that it is correct is calculated based on the number of times the Wikipedia article for that sense is used as a link destination in Wikipedia (its prior probability), the semantic relatedness of the sense with other article references in the neighboring context, and a measure of the overall quality of the context. This calculation uses a machine learning model whose parameters have been determined in a training phase that uses manually assigned links in Wikipedia as ground truth.

The next task is to identify which terms in the document should be linked. Terms are considered linkworthy if they are used as a link in many Wikipedia articles; are closely related to the main topic of the document; are mentioned numerous times in the text; are located near the beginning of the document; and there is a high confidence that the term has been correctly disambiguated. Again a machine learning model is used that has been trained on the existing links in Wikipedia.

Once the key terms in the text have been identified and linked to corresponding Wikipedia articles, the annotated text is passed back to the Realistic Books software. All links are converted to HTML hyperlink tags. For each key term, the first paragraph and accompanying image is retrieved from the Wikipedia article and set as the tooltip text for the link to provide readers with immediate access to definitions.

All terms in the book that point to the same Wikipedia article are inserted into a synonyms table, along with redirect terms for the article that are defined in Wikipedia. Then the back-of-the-book index is generated and neighboring entries conflated as illustrated above for *Obama, Pres. Obama*. (Capitalization of entries in the index follows the capitalization used in the book.) Finally, the original HTML and metadata tags are inserted into a fresh copy of the book content file.

7. System responsiveness

Responsiveness is the speed with which a system responds to a user's action, and significantly affects user satisfaction, engagement and reading performance (Hansen & Haas, 1988). To investigate the responsiveness of Realistic Books, we recorded the average time required to load a document, turn a page and perform a text search on 25 books with 500–2500 pages (with no blank pages), and 1100 words per page, at two different screen sizes (800×600 and 1440×900). To display an open book with enough space beneath for the bottom of the stack of pages and the book cover (see Figure 2), and enough space above to turn a page, requires a screen that is somewhat larger than the double-page spread itself, and these two screen sizes comfortably accommodate books with page sizes of approximately 330×460 and 495×700 respectively—although we experimented with a wider variety of A4 sizes (from 150×210 to 2400×4000 pixels) because readers can zoom in or out.

All timings were measured on a Windows computer with a 2.8 GHz Intel Pentium 4 processor, and an Apple laptop with a 2.2 GHz Intel Core 2 Duo processor; the figures for both were the same to within 10%. To eliminate the influence of Internet connection speed, documents were viewed offline. The page turning speed was set to *medium*, which means that there are 11 poses during each turn (including the initial and final ones).

7.1. Opening the book

The Realistic Books system loads pages on a demand basis. When a book is opened at a new page—at a random position, rather than as the result of a page-turning action—six pages are cached: the current double-page spread and two pages before and after (in preparation for backward and forward page turns). When a book is displayed initially, only the front cover is shown, and the following two pages are loaded in preparation for a forward page turn. However, first the metadata for the entire book is read in, which incurs a delay that depends on the length of the book and the amount of metadata.

For typical books, the start-up delay ranges from 2 secs for 500-page books to 7 secs for 2500-page ones. (This figure compares favorably with that for the Adobe Reader when reading linearized PDF

documents, whose measured start-up time for these books is about 3.5 and 8.5 secs respectively; when a browser plug-in is used instead of the Adobe Reader a further 1 sec is required to load the plug-in first.) The time to load a page is, in practice, independent of page size, because it is dominated by the need to render the text on the page: it was measured at 10 msec/page. Thus it takes around 60 msec to load the 6 pages that are cached when the book is opened at a random position.

7.2. Turning a page

Turning a page of a Realistic Book requires computing the appearance of the four pages involved at each intermediate position and displaying the result; then updating all state variables and loading the next two pages in preparation for the next page-turn. Excluding the actual display operation, the time required is independent of the screen resolution, the size of the page and the total number of pages: it was measured at about 70 msec plus the two page loads (10 msec/page). The time taken to draw the book at each position depends on the page size, and ranges from 60 msec for small pages (150×210 pixels) to 140 msec for larger ones (600×850 pixels). The time remains constant thereafter as the pages become larger, because Flash only masks those object parts that lie within the user's viewing area, which now becomes limited by screen size. Thus with 10 intermediate poses (excluding the initial one but including the final one), it takes slightly over 0.5 sec to completely turn a small page and 1.5 sec to turn a large one—no matter how large—on a 1440×900 pixel screen.

7.3. Text search

As noted earlier, Realistic Books implement a form of incremental full-text search that performs partial matching of words or phrases in the entire book, as well as matching against a separate table that contains keyphrases that have been identified in the book, and their synonyms. Full-text search uses a trie data structure that represents the entire book and gives the numbers of all pages that contain any character string in it, along with the character offsets within that page of where the string appears. The keyphrase/synonym table gives page numbers and character offsets of all the terms it contains. In either case, the time required for searching is independent of the size of the book.

First, the current string that has been typed into the search box is sought in the keyphrase/synonym table—this takes about 1 msec. If the search is successful, a list of pages and character offsets is returned. Otherwise, the trie data structure is consulted to locate all pages that contain the string. In either case, the relevant pages are bookmarked. These operations take about 1 msec in the first case and 12 msec in the second, even for a very common search string (e.g., the letter *e*) and a large book. Then all occurrences of the search string in the current double-page spread are highlighted, which takes a total of about 20 msec—again for a very common search string. At this point, the user sees the result of the search. Finally, all occurrences of the search string on the remaining four pages that have been loaded in preparation for a forward or backward page-turn are highlighted, taking a further 40 msec. As the user appends more characters to their query, less time is required to perform all these operations. Thus the entire search operation is accomplished within 72 msec in the worst case.

8. Usability studies

One way of determining whether electronic books actually improve people's reading is to compare their performance with commonly used document representations, and with physical books, on tasks that represent likely situations in which the system will be used (Gould, 1988). The typographical features, logical structure and page layout of the electronic text should all be equivalent to the paper material. Moreover, evaluations should be designed not only to show whether the new system improves upon the others, but to identify which of its properties are responsible for any improvement.

Four separate evaluations were conducted to verify that the design interface and reader services implemented in Realistic Books can be understood and easily used by readers, and to collect information for

improving the current realization. We compared Realistic Books with a web browser, the Adobe PDF Reader, and physical books, in terms of how people browse (Section 8.3), find information using plain search (Section 8.4), find information in documents that have been semantically enriched as described in Section 6 (Section 8.5), and personalize documents (Section 8.6). Evaluating electronic book readers such as Amazon's Kindle and Sony's Digital Book lie beyond the scope of this research. (However, this may not be a significant shortcoming, because their display resembles the PDF format—without its scrolling capability.) This section describes the experimental design and criteria employed in the user studies, and then reports the results and conclusions drawn from each evaluation.

8.1. Design

Procedures. Subjects came individually to each study and completed a profiling questionnaire that recorded age range, gender, language proficiency, and their experience with computers, web browsers and Adobe Reader. To ensure that they understood how to use all the functions that they might need, they were asked to familiarize themselves with each document format and were then explicitly shown any functions that they did not appear to know. Once ready, they were given lists of tasks and asked to perform them.

A task based scheme was developed for a comparative user-centered evaluation between Realistic Books and equivalent documents in HTML, PDF and printed book form. For each task, the the success score and time taken to complete the task were recorded. To ensure that they felt no pressure and worked at their own pace, subjects did not know that any timings were being recorded. They were encouraged to voice their thoughts and feelings as they worked, and the functions that they used during the task were logged. During the experiment, the evaluator made copious notes of each user's actions.

A "repeated measures" usability design was employed, in which the same participants were tested under every possible condition (Dix et al., 2003). For example, half the subjects performed Task A with physical books and Task B with Realistic Books; the other half performed Task A with Realistic Books and Task B with physical books—where Tasks A and B were similar but not identical. This procedure eliminates experimental bias caused by individual differences between participants, and minimizes the required number of participants. However, performance on a later task may be affected by the experience of having performed other tasks (Kinnear & Gray, 2006). To balance this effect across the conditions, the various document types were presented to subjects in a different order.

Materials. Documents in all formats except HTML were paginated in exactly the same way, and for PDF the beginning of each section was bookmarked. The HTML books were presented according to the Wikipedia article presentation style. The standard Adobe Reader was used to view PDF documents. In cases where a physical book was not available, one was created by printing the document double-sided in color, and stapling it down the left-hand side. To avoid delays in loading or refreshing, all documents were served off-line, having been fully loaded before the experiment began.

A total of seven books were used for the evaluation studies: three for comparing electronic document formats and a further two for comparing with physical books, both in Section 8.3; one for the searching studies in Sections 8.4 and 8.5; and another for the personalization study in Section 8.6. These are briefly described in the sections below, along with the tasks themselves; an Appendix gives further details.

Participants. In all experiments there were at least eight participants for each treatment condition. Obviously readers vary widely in how they understand a text and the way they process it. To help mitigate the effect of variability, high school and university students aged 15–40 and from a variety of disciplines were recruited. In order to bias the experiment *against* the Realistic Books format, participants were chosen who had used computers extensively for several years and were familiar with web browsers and PDF readers.

Participants' views and subjective opinions on their experience were gathered through the use of questionnaires. They had to choose which formats were relevant and useful for them, easier to navigate or annotate, easier to locate information, more pleasant and engaging to use, and finally which format they preferred overall—with the possibility of choosing more than one format. They were then asked to explain their choices and state which features they liked and disliked, and suggest improvements.

8.2. Criteria

Reading behavior and performance can be assessed in terms of outcome and process measures (O'Hara & Sellen, 1997). The former focus on what readers obtain from the text, how usable the system is, and whether it provides functions that allow people to accomplish their tasks. The latter assess how readers engage with the material, and are obtained by observing participants' reading behavior.

ISO 9241–11 (1998) defines effectiveness, efficiency and satisfaction as the three main criteria for outcome measures. Quesenbery (2001) and Shneiderman (1998) further expand the ISO definition by explicitly specifying items that need to be evaluated for each criterion. *Effectiveness* can be measured by the number of tasks a user completes successfully. *Efficiency* concerns the time required by subjects to complete a task. *Satisfaction* refers to participants' opinions in terms of pleasantness, sense of engagement, ease of use, ease of learning, preference and utility.

Based on these criteria, four experiments were designed to determine whether users of Realistic Books would satisfy the following hypotheses:

Browsing (see Section 8.3)

- H1.1 Find answers, through browsing, faster than with other document formats
- H1.2 Make fewer task errors than with other formats
- H1.3 Report higher satisfaction than with other formats

Plain text searching (see Section 8.4)

- H2.1 Find answers, through searching, faster than with other document formats
- H2.2 Make fewer task errors than with other formats
- H2.3 Report higher satisfaction than with other formats
- H2.4 Use the most effective search tools to answer questions, regardless of document format, i.e., back-of-the-book index for index questions and search tools for full-text search questions

Enriched document representation (see Section 8.5)

- H3.1 Find answers faster in enriched documents
- H3.2 Make fewer task errors
- H3.3 Report higher satisfaction

Personalization (see Section 8.6)

- H4.1 Report higher satisfaction than with other document formats.

8.3. Browsing

Readers often scan pages casually in the hope of finding something of interest, using navigation mechanisms such as hyperlinks, document structure cues and navigation buttons while browsing visually through the text. To quantitatively evaluate the browsing experience in Realistic Books against conventional formats (web browser and PDF reader), Spool et al.'s (1997) scavenger hunt approach was adopted; then a follow-up comparison was conducted with physical books. The task involves browsing through test material in search of specific information—without using search commands.

In the first study participants were asked to read three simple books: a handbook (*Raising Ducks: how to begin*), a history book (*History of the United States told in one-syllable words*) and a children's book (*Alice's Adventures in Wonderland*), each with

80–140 pages divided into 8–18 chapters. For the follow-up comparison with printed books, subjects were asked to read two biographies from the Internet Archive (*Pavlova: A biography*, 84 pages, and *Princess Mary: A biography*, 254 pages), both with 10 chapters.

While viewing the book, participants answered six multiple-choice questions. Three were aimed at determining whether they understood what it was about, and could be answered from the front matter and introductory section. The remaining three were designed to assess critical reading activity and required participants to search through the text.They were unlikely to know the answers already, and often needed to browse more than one section to find the information and draw a conclusion by relating several pieces of supporting evidence. The Appendix summarizes the questions used for this and other tasks.

Efficiency. A one-way analysis of variance of the time taken to answer the questions in the first browsing study shows that the differences due to document format were statistically significant at the 5% level ($F_{2,21} = 4.69$). Further t-test analysis shows that the difference between Realistic Books and HTML was significant at the 5% level ($t_{21} = 3.09$), and at the 10% level between Realistic Books and PDF ($t_{21} = 1.54$). Because a one-way analysis of variance for the second browsing study showed no significant difference between Realistic Books and PDF ($F_{1,14} = 1.59$), H1.1 (*speed*) was only partially upheld.

Effectiveness. H1.2 (*accuracy*) was not upheld. A one-way analysis of variance of the number of correctly-answered questions detected no significant differences due to document format (first browsing study: $F_{2,21} = 0.41$, second study: $F_{1,14} = 0.04$). We conclude that although document format does not affect the reader's level of comprehension, it does affect the time required to process text.

Reading experience. Almost all readers (80%) commented that while web browsers suffice for short documents, they prefer paginated displays for longer ones. Without the page edges that Realistic Books and physical books show, 85% of the subjects felt lost in the documents. They did not know where they were; sometimes did not notice when they moved to the next section; and found it difficult to return to a previously encountered passage. All found Realistic Books to be as useful, engaging and easy to use as physical books, and preferred them over web browsers and the PDF reader. Thus H1.3 (*satisfaction*) was partially upheld.

Liesaputra & Witten (2008) give more information on the design and results of the browsing studies.

8.4. Searching

Subject indexes are an important access tool for physical books, whose utility frequently depends on the quality of the index. But as text migrates to electronic formats that are searchable by computer, it is easy to assume that indexes become superfluous. However, readers who do not know exactly what they seek or how it might be expressed may issue search queries that match many irrelevant passages—or none at all. Advanced search functions, which could be used to improve search results, often make people anxious because they lack experience with them.

To investigate search behavior and performance with Realistic Books, they were evaluated against the web browser, PDF reader, and printed books. Participants read a university level text called *Data Mining: Practical Machine Learning Tools and Techniques*, which has a professionally produced subject index whose terms accurately represent the contents. To reduce the search space, only the first three chapters were used: the table of contents and back-of-the-book index were modified to remove references to other chapters. The Realistic Books *include synonyms* search option was turned off to ensure that automatic term aliasing did not affect the result.

There were four questions, which participants undertook in the same order, summarized in the Appendix. Two were intended to be answered by text search and two by consulting the subject index, and they were designed so that it was less efficient to answer

text-search questions using the back-of-the-book index than by using the full-text search function, and similarly for the subject index questions. The intention was to force participants to use the search tool and the back-of-the-book index at least once while working with each document format. We wished to determine whether readers preferred to use the full-text search tool, the back-of-the-book index or both.

Efficiency. A one-way analysis of variance of the time taken to answer the questions shows that differences due to format were statistically significant at the 1% level ($F_{3,28} = 5.09$). However, t-tests show no significant difference between the web browser and PDF reader for any task ($t_{28} = 0.08$), while users found answers significantly faster (at the 10% level) using the web browser and PDF reader than in physical books (browser: $t_{28} = 1.39$, PDF: $t_{28} = 1.31$). This upholds H2.1 (*speed*): the improvement from Realistic Books to physical books was statistically significant at the 1% level, and their improvement over the web browser and PDF reader was significant at the 5% level (physical: $t_{28} = 3.06$, browser: $t_{28} = 1.71$, PDF: $t_{28} = 1.75$).

Effectiveness. A one-way analysis of variance found no significant differences between document formats in the number of task errors made ($F_{3,28} = 0.16$). Thus, as with the browsing studies, H2.2 (*accuracy*) was not upheld: accuracy was not influenced by document format.

Search experience. Physical books were preferred for reading activities, but not for information seeking. Participants thought that Realistic Books combine a good reading environment with a good search experience, and rated them the most useful, engaging and easy to use of all, upholding H2.3 (*satisfaction*). HTML was the next preferred format. However, subjects were not satisfied with simple string-match search. They wanted typing errors to be detected and synonyms to be incorporated, to avoid having to consult the back-of-the-book index.

Search tool preference. All participants relied on full-text search to find answers in all electronic formats, and the back-of-the-book index was considered as a secondary strategy only in Realistic Books. Thus H2.4 (*appropriate choice of search tool*) was only partially upheld. When searching for information in HTML or PDF books, nearly everyone neglected the back-of-the-book index, even though they had been told about it before commencement and it was listed in the table of contents. Readers only stumbled upon the index when it contained a search term that appeared nowhere else—usually a synonym or bridging word for the actual terminology in the text.

Liesaputra (2010) provides more information on the design and results of this user study.

8.5. Enriched document representation

Electronic documents should provide ways for users to quickly spot and understand keywords, and rapidly locate a topic and all related pages (Poynter Institute, 2000; Goldsborough, 2000; Berkenkotter & Huckin, 1993; Crane, 2002; Dillon et al., 1989; Murray, 2002). If a back-of-the-book index is not available, Crestani & Ntioudis (2001) suggest that an electronic book application should automatically create one. Following their suggestions, Realistic Books support the above activities through automatically generated definition links, synonyms, and the back-of-the-book index, as described in Section 6. This experiment evaluates the utility of text with and without these facilities in order to assess the effect of the enriched document representation.

Participants were asked to read the abridged *Data Mining* book in Realistic Books format. The subject index pages were removed to ensure that participants could not locate better query terms from the manually created index. Three versions of the book were made, the first with navigational hyperlinks in the table of contents, the second with navigational hyperlinks and a back-of-the-book index, which also includes synonym search, and the third with all enrichment facilities, including definition popups and Wikipedia links in the main text.

There were three sets of tasks, which participants performed in the same order, and each task asked three multiple-choice questions. Although none of the terms used in the questions appears explicitly

in the book, a superficial understanding of the text would be sufficient to guess an appropriate query term. To provide a balanced assessment, the questions had varying degrees of difficulty. The first was a general question designed to be easy to answer, and the term it mentioned was present in the synonym list. The other two assessed analytical reading activity and required users to identify all the pages that mention a specified topic. In one case the term specified in the question was present in the synonym list; in the other it was not.

Efficiency. A one-way analysis of variance of the time taken to answer questions showed that the document enrichment improves performance for all tasks (statistically significant at the 1% level, $F_{2,21}$ = 7.60), which upholds H3.1 (*speed*). The improvement from the first version (hyperlinked table of contents only) to either of the other two was statistically significant (1% level, second version: t_{21} = 2.95, third version: t_{21} = 3.68), but there was no significant difference between the second and third versions, that is, the inclusion of definitional popups and Wikipedia links in the main text made no significant difference (t_{21} = 0.73).

Effectiveness. Overall, participants made fewer errors when answering with the enriched representation, and obtained better scores with definition links present than without. However, a one-way analysis of variance of the number of correct answers found no significant difference ($F_{2,21}$ = 0.12), so H3.2 (*accuracy*) was not upheld.

Search experience. All participants preferred the enriched document representation, upholding H3.3 (*satisfaction*). They reported that the search function was more flexible and error tolerant because it automatically locates terms related to the query term, and in some cases handles misspelling. Because searching for a given term automatically bookmarks pages containing its synonyms as well, and the synonym pops up when the bookmark is moused over, readers gain a better understanding of the text.

Although all subjects relied on full text search to find answers, 60% utilized the back-of-the-book index as a secondary search strategy, scanning the index to help them obtain better search terms. However, most (55%) were disappointed with the quality of the index, feeling that it contained too many terms considering the brevity of the (abridged) book.

Most participants (70%) preferred the text with definition links because they felt that they could scan passages more quickly, understand the text better and formulate more suitable query terms. Few bothered to click on a definition link to obtain a detailed explanation of the term in Wikipedia, reporting that the brief pop-up definitions were enough for them—and those 15% who did follow the links remarked that once in Wikipedia they were tempted to browse unrelated material and needed a conscious effort to focus attention back to the set task.

8.6. Personalization

People who personalize their documents usually do so to mark information as interesting. They presumably believe that adding notes or bookmarks will yield benefits that exceed the manipulation costs of creating and managing them. Because users find it easy to personalize paper documents, annotating and bookmarking are an integral part of traditional reading practices (Adler et al., 1998). In contrast, online readers found personalization to be too disruptive and to require more effort than a pencil or highlighter (O'Hara & Sellen, 1997; Marshall, 1997). Therefore they often opt not to personalize their electronic documents, even when they can foresee a long-term advantage.

This experiment investigated users' opinions on the costs of creating and managing personalizations in Realistic Books, physical books, and PDF files (HTML was not used because the annotation plugins that support personalization offer similar facilities to PDF).

Subjects were asked to read Chapters 6, 7 and 11 of a university level text (*Search user interfaces*), chosen because they are small (10 to 12 pages) and can be understood without background knowledge of the subject matter. The table of contents was modified by removing all references to other parts of the book;

there was no index. Participants were provided with Post-It notes, bookmarks, pencils, pens and highlighters in different colors and sizes. They were asked to undertake a series of tasks that explored all the basic personalization functions provided by each document format, such as highlighting passages, creating annotations and bookmarking pages.

Unlike previous experiments, this study focused on participants' experience and strategies, not their performance. There was no significant difference between the formats in the average number of personalizations created and the quality of the annotations. It was observed that subjects tried to transfer their behavior from paper to the computer. They each utilized at most two personalization tools; highlighting and underlining were preferred to annotating.

In general, subjects felt that all the personalization functions were useful and easy to use. A few (20%) said they preferred to type rather than write their annotations because sometimes they cannot understand their own handwriting when they revisit a note. Participants preferred to personalize Realistic Books over PDF and physical books, upholding H4.1 (*satisfaction*).

Subjects noted that Realistic Books and physical books provide a smoother transition from reading to personalizing than the PDF reader, where they had to explicitly switch between different tools. Many (60%) complained that they did not know whether a PDF page was bookmarked without consulting the bookmark view. Adobe Reader supports many personalization functions, but almost all subjects (90%) were confused as to what the icons represent.

Many users preferred Realistic Books over physical ones because they could move, edit and search the annotations. They felt that they could write as much as they wanted without worrying about the space on the page. However, 40% preferred physical books because they are more familiar and feel more comfortable—particularly with regard to the fluidity of scribbling—and because physical books can be read while not at a workstation.

8.7. Findings

Current digital document systems replace or remove some of the orientation cues available in paper books. However, our user studies show that when these cues are removed participants feel lost, disoriented and frustrated. To identify which printed and electronic document features are essential, we recorded the frequency of each feature used during the studies. This section describes the orientation cues and interactions that participants considered important and used most often.

Hyperlinked table of contents and back-of-the-book index. Both function as maps that inform readers about the depth of the information space, text organisation, and possible ways to reach specific information. Few electronic documents have a back-of-the-book index. Subjects appreciate the fact that Realistic Books could automatically generate back-of-the-book indexes and support full-text search over synonyms.

Card et al. (2004a) experimented with reorganising back-of-the-book indexes, and found that readers locate information quicker and more accurately if they are only shown index entries that are relevant to the search query. However, our observations with people's searching behaviour found that the back-of-the-book index was considered as a secondary strategy in Realistic Books: they looked at it only when they could not think of a good search term; then they browsed through the index to find terms for their query. In the future, it would be interesting to see whether adding the reorganised back-of-the-book index feature into Realistic Books could further improve searching experience and performance.

Bookmarks make a page accessible from any other page. Because they are always visible on the book's page edges, readers can use them to indicate important pages and estimate the size of sections and the reader's position in them. Bookmarks are also used as a point of reference for backtracking. Having referred to the index or the table of contents page, readers can use the "previously visited page" bookmark that Realistic Books show to return to their original location and continue.

Very few printed books, except for dictionaries, bibles and encyclopaedias, explicitly mark chapters in a way that can be seen even when the book is not open at that page. However, participants liked the fact that Realistic Books can automatically and dynamically bookmark chapter and section openings, the table of contents and back-of-the-book index, pages with images and annotations, and all pages related to a specific topic. Taking advantage of the book's spatial layout, this feature makes those special pages accessible from any page and provides users with a visual indication of the document's logical structure.

While performing text search in both Realistic Books and HTML, subjects routinely utilized the bookmarks that mark the location of matched search terms to navigate between the search results. Although they still preferred the bookmarks in page edges over the coloured markings in the scroll bar, this observation indicates that reading experience with HTML and PDF documents may be improved by automatically placing colored markings in the scroll bar to indicate chapter and section openings, parts of documents with images or annotations, and previously visited text. In the future, we plan to compare the performance of Realistic Books with that of new HTML and PDF document representations that have been augmented with colored markings in the scroll bar and given the ability to collapse parts of the document.

Double-page spread, full-page view, typographic cues, illustrations and *annotations* provide readers with a broader context of the article and how its information is organised. As a result, unusual text layout, illustrations, comments, scribbles, highlighted passages, bookmarks and section titles can all be utilised as visual landmarks. With a partial-page view or a single-page spread, participants cannot immediately tell whether one page has been typeset differently than the others because this view does not provide such a broader context of the text. Without any annotations or illustrations, participants felt that the texts always look the same.

Pagination helps break the flow of information in a long document. Page size, running header and page number define the physical boundary of the text in a page. Although page numbers and running headers can be used as landmarks that inform users of their current position, few people notice or remember this information. Instead, readers utilise any visual landmark surrounding the text, the appearance of the stack of page edges or the position of the scroll bar slider to return to an approximate location.

Subjects commented that they would prefer to alter the font size to make the text more readable as opposed to modifying the document's zoom factor and not being able to view the content of the whole page. In the future, we are planning to incorporate this feature into Realistic Books.

The *stack of page edges* on either side of the book suggest the size of the document and the reader's current position. With HTML and PDF, participants did not consider the size of the scroll bar's slider as a good indication of the document's length. They tended to scroll to the end of the document to judge its length and utilise visual landmarks to determine their position. Furthermore, they found it hard to return to a known location with a scroll bar because a slight change in slider position can significantly change the part of the document shown.

Fluid interaction between browsing, searching and personalising helps readers find information faster and gives them a more pleasant experience. People often look back and re-read some text to get a better understanding of the passage they are currently reading. They also look ahead in the document to determine the length of the current section and to decide whether it is worth continuing. When they are critically reading a document, subjects usually annotate it to mark important passages or insert their comments. These activities are all intertwined. Unlike in physical and simulated books, in HTML and PDF these interactions are interruptive rather than being unselfconsciously integrated into the flow of regular activities. Hence people tend to shy away from them.

Interactive page turning simulation makes document more engaging, ensures that readers never lose contact with the text, indicates the direction of travel, and helps people remember how many pages they

have read. It also makes it easier to return to known information, because readers remember how many pages to turn.

The act of riffling through pages can be preserved simply by displaying a list of the page thumbnails. People can browse through the list looking for pages with distinctive visual structure while still viewing the page they are currently reading, and find interesting pages by accident. However, the physical act of riffling pages, the backtracking mechanism and the ability to change the document's magnification level were rarely used during the studies and are therefore the least important features to be incorporated into electronic book applications.

Key phrases in Realistic Books can be automatically identified and cross-referenced to Wikipedia articles, providing readers with another form of visual landmark and the ability to search over synonyms. Participants appreciated the fact that they can just hover over a key phrase to obtain its definition: they do not have to shift their focus away from the text to the glossary page or to an external dictionary to obtain the term's definition.

9. Conclusion

People rarely read books linearly; they skip around. They may search for things they have seen before and for things they have not, browse for new information, and assess a book's quantity and coverage of information without regard to the details (Canter et al., 1985). They rely on orientation cues provided by the document interface to know their current position, the size of the document, where to go next and how to visit parts that are not in view. Although much of a book's physical design is devoted to orientation cues, most electronic displays dispense with them. Some researchers believe that presenting double-page spreads or displaying the stack of page edges on either side of a book wastes valuable screen space (Nielsen, 1998b; Rowland et al., 1997).

To determine whether it is worth simulating the look and feel of physical books in electronic environments, the Realistic Books has been developed. This paper describes the facilities that it provides and how it is implemented; investigates user performance and strategy while browsing, searching and personalizing; and compares this model with web browsers, PDF readers, and physical books. Subjects preferred Realistic Books over the other formats, and were found to complete tasks significantly faster. Readers can navigate and annotate Realistic Books as easily as printed books while retaining the advantages of an electronic environment—such as searching, editing, accessing multimedia, and automatic semantic enrichment. In line with earlier studies, no significant differences were found between the formats in terms of accuracy and user comprehension (Kak, 1981; Muter et al., 1982; Cushman, 1986; Dillon et al., 1988).

People tend to employ the same tools and strategies when reading electronic text as they do for physical books. This paper has shown how to simulate documents using the book metaphor. Page turning and page edges help subjects browse around the material. The personalization tools implemented in Realistic Books provide a relatively seamless integration between reading and annotating. The automatically generated definition links, synonyms and back-of-the-book indexes help users understand the content better and find related information faster. Most importantly, Realistic Books are easy to learn and use.

Several Realistic Books are available at `www.nzdl.org/books` for readers to peruse themselves. Our software to make books from HTML and PDF files is open source and available at the same location.

Acknowledgments

We acknowledge the entire New Zealand Digital Library Project team for their unstinting work in providing an environment that makes this kind of research meaningful—and fun. This research is funded in part by Google.

Adler, A., Gujar, A., Harrison, B. L., O'Hara, K., & Sellen, A. (1998). A diary study of work-related reading: Design implications for digital reading devices. *Computer-human interaction*, (pp. 241–248).

Adler, M. J., & van Doren, C. (1972). *How to read a book*. New York: Simon and Schuster.

Ajmera, R., Borgaonkar, S., Madhvanath, S., Amulya, J., & Kozhissery, R. (2009). Papyrus - design and evaluation of a dual display notebook PC form factor enabling writing/annotation and parallel tasking. *HP Lab*. Available at `http://www.hpl.hp.com/india/documents/papers/HPL-2009-81.pdf`.

Apple (2010). iBooks: A novel way to buy and read books. Available at `http://www.apple.com/ipad/features/ibooks.html`.

Baines, J. (1983). Literacy and ancient Egyptian society. *Man*, *18*, 572–599.

Barker, P., Richard, S., & Benest, I. (1994). Human computer interface design for electronic books. *Proceedings of the* 18th *international online information meeting*, (pp. 213–225).

Benest, I. D. (1990). A hypertext system with controlled hype. *Hypertext: State of the art*, (pp. 52–63).

Berkenkotter, C., & Huckin, T. (1993). Rethinking genre from a sociocognitive perspective. *Written communication*, *10*, 475–509.

Bjork, S. (2001). The ScrollSearcher technique: using scrollbars to explore search results. *INTERACT'01*, (pp. 745–746).

British Library (2006). Turning the Pages. Available at `http://www.bl.uk/onlinegallery/ttp/ttpbooks.html`.

Buchanan, G., & Pearson, J. (2008). Improving placeholders in digital documents. *ECDL 2008*, (pp. 1–12).

Burbules, N. C. (1998). Rhetorics of the Web: Hyperreading and critical literacy. *Page to screen: Taking literacy into the electronic era*, (pp. 102–122).

Canter, D., Rivers, R., & Storrs, G. (1985). Characterizing user navigation through complex data structures. *Behaviour and information technology*, *4*, 93–102.

Card, S. K., Hong, L., Mackinlay, J. D., & Chi, E. H. (2004a). 3Book: A 3D electronic smart book. *Proceedings of the working conference on advanced visual interfaces*, (pp. 303–307).

Card, S. K., Hong, L., Mackinlay, J. D., & Chi, E. H. (2004b). 3Book: A scalable 3D virtual book. *Extended abstracts of the ACM international conference on human factors in computing systems (CHI 2004)*, (pp. 1095–1098).

Card, S. K., Robertson, G. G., & York, W. (1996). The WebBook and the WebForager: An information workspace for the World-Wide Web. *Proceedings of the SIGCHI conference on human factors in computing systems: Common ground*, (pp. 111–117).

Chen, C. (2000). Individual differences in a spatial-semantic virtual environment. *Journal of the American Society for Information Science*, *51*, 529–542.

Chen, N., Guimbretiere, F., Dixon, M., Lewis, C., & Agrawala, M. (2008). Navigation techniques for dual-display e-book readers. *Proceedings of CHI'08*, (pp. 1779–1788).

Choi, K. (2010). Significance of trade. Available at `http://www.econ.iastate.edu/classes/econ355/Choi/sig.htm`.

Chu, Y.-C., Bainbridge, D., Jones, M., & Witten, I. H. (2004). Realistic books: A bizarre homage to an obsolete medium? *JCDL '04: Proceedings of the* 4th *ACM/IEEE-CS joint conference on digital libraries*, (pp. 78–86).

Chu, Y.-C., Witten, I. H., Lobb, R., & Bainbridge, D. (2003). How to turn the page. *JCDL '03: Proceedings of the* 3rd *ACM/IEEE-CS joint conference on digital libraries*, (pp. 186–188).

Cipolla, C. M. (1969). *Literacy and development in the West*. Harmondsworth: Penguin.

Cockburn, A. (2004), .

Cockburn, A., & Greenberg, S. (1999). Issues of page representation and organization in Web browser's revisitation tools. *Australian journal of information systems*, *7*, 120–127.

Cockburn, A., & McKenzie, B. (2001). 3D or not 3D? Evaluating the effect of the third dimension in a document management system. *Proceedings of CHI 2001 on human factors in computing system*, (pp. 434–441).

Cockburn, A., & McKenzie, B. (2002). Evaluating the effectiveness of spatial memory in 2D and 3D physical and virtual environments. *Proceedings of CHI 2002 on human factors in computing system*, (pp. 203–210).

Conklin, J. (1987). Hypertext: An introduction and survey. *IEEE Computer*, (pp. 17–41).

Crane, G. (2002). Cultural heritage digital libraries: Needs and components. *Proceedings of* 6th *European conference on research and advanced technology for digital libraries*, *2458*, 626–637.

Crawford, W. (1998). Paper persists: Why physical library collections still matter. *Online magazine*, (pp. 42–48).

Crestani, F., & Ntioudis, S. P. (2001). User centred evaluation of an automatically constructed Hyper-TextBook. *Journal of educational multimedia and hypermedia*, (pp. 3–19).

Cushman, W. H. (1986). Reading from microfiche, VDT and the printed page: Subjective fatigue and performance. *Human factors*, *28*, 63–73.

Czerwinski, M., vanDantzich, M., Robertson, G., & Hoffman, H. (1999). The contribution of thumbnail image, mouse-over text and spatial location memory to web page retrieval in 3D. *Proceedings of interact 1999*, (pp. 163–170).

Dillon, A. (1992). Reading from paper versus screens: A critical review of the empirical literature. *Ergonomics*, *35*, 1297–1326.

Dillon, A. (2000). Spatial-semantics: How users derive shape from information space. *Journal of the American Society for Information Science*, *51*, 521–528.

Dillon, A., McKnight, C., & Richardson, J. (1988). Reading from paper versus reading from screens. *The computer journal*, *31*, 457–464.

Dillon, A., Richardson, J., & McKnight, C. (1989). Human factors of journal usage and design of electronic texts. *Interacting with computers*, *1*, 183–189.

Diringer, D. (1953). *The hand-produced book*. New York: Philosophical Library.

Dix, A., Finlay, J., Abowd, G. D., & Beale, R. (2003). *Human-computer interaction*. Harlow: Prentice Hall.

Drucker, J. (2006). The virtual codex from page space to e-space. Available at http://www.philobiblon.com/drucker/.

Dumas, C., Jourdan, F., & Plenacoste, P. (2002). Libraries: Comparisons between the real and the virtual in 3D, 2D zoomable and 2D arborescent. *In JCDL workshop usability of digital libraries*, (pp. 11–14).

Dunn, J. (2006). Discovering writing. Available at http://www.touregypt.net/historicalessays/discwriting.htm.

Easley, M. F. (2008). Dead sea scrolls coming to N.C. museum of natural sciences. Available at http://www.enr.state.nc.us/upclose/pages/bulletinboard27.html.

Edwards, D., & Hardman, L. (1989). Lost in hyperspace: Cognitive mapping and navigation in a hypertext environment. *Hypertext: theory into practice*, (pp. 90–106).

Fuller, D. (2006). Honors scholars get a world lesson exploring books in human culture. Available at http://www.uc.edu/news/NR.asp?id=3707.

Furnas, G., Landauer, T., Gomez, L., & Dumais, S. (1987). The vocabulary problem in human-system communication. *Communications of the ACM*, *30*, 984–971.

Gaur, A. (1984). *A history of writing*. London: British Library.

Genuth, I. (2007). E-paper history: an interview with Nick Sheridon, father of e-paper. Available at http://thefutureofthings.com/articles/1000/the-future-of-electronic-paper.html.

Goldsborough, R. (2000). Text demands respect on the Web. *Advertising age*, *71*, 44.

Gomez, L., C.C., L., & T.K., L. (1990). All the right words: Finding what you want as a function of richness of indexing vocabulary. *Journal of the American society for information science*, *41*, 547–559.

Good, M., Whiteside, J., Wixon, D., & Jones, S. (1984). Building a user-derived interface. *Communications of the ACM*, *27*, 1032–1043.

Gordon, S., Gustavel, J., Moore, J., & Hankey, J. (1988). The effects of hypertext on reader knowledge representation. *Proceedings of the human factors society*, (pp. 296–300).

Gorman, M., & Crawford, W. (1995). *Future libraries: Dreams, madness and reality*. Chicago, IL: American Library Association.

Gould, J. D. (1988). How to design usable systems. *Handbook of human-computer interaction*, (pp. 757–789).

Gould, J. D., Alfaro, L., Barnes, V. M., Finn, R., Grischkowsky, N., & Minuto, A. (1987). Reading is slower from CRT displays than from paper: Attempts to isolate a single-variable explanation. *Human factors*, *29*, 269–299.

Hansen, W. J., & Haas, C. (1988). Reading and writing with computers: A framework for explaining differences in performance. *Communication of ACM*, *37*, 1080–1089.

Harris, W. V. (1989). *Ancient literacy*. Cambridge, MA: Harvard University Press.

Henke, H. (1998). Are electrons better than papyrus? (or can Adobe Acrobat Reader files replace hardcopy? *ACM SIGDOC conference on scaling the heights*, (pp. 29–37).

Hill, W. C., Hollan, J. D., Wroblewski, D., & McCandless, T. (1992). Edit wear and read wear. *Proceedings of the SIGCHI'92 conference on human factors in computing systems*, (pp. 3–9).

Hinckley, K., Dixon, M., Sarin, R., Guimbretiere, F., & Balakrishnan, R. (2009). Codex: a dual screen tablet computer. *Proceedings of CHI'09*, (pp. 1933–1942).

Hinckley, K., Yatani, K., Pahud, M., Coddington, N., Rodenhouse, J., Wilson, A., Benko, H., & Buxton, B. (2010). Pen + touch = new tools. *UIST'10*, (pp. 27–36).

Internet Archive (2007). BookReader. Available at http://www.OpenLibrary.org/dev/docs/bookreader.

ISO 9241–11 (1998). Ergonomic requirements for office work with visual display terminals (VDTs)—Part 11: Guidance on usability, .

Kahan, M. R., J. ad Koivunen (2001). Annotea: An open RDF infrastructure for shared Web annotations. *WWW10 international conference*, (pp. 623–632).

Kak, A. V. (1981). Relationships between readability of printed and CRT-displayed text. *Proceedings of human factors society annual meeting*, (pp. 137–140).

Kearsley, G. (1988). *Online help: Systems design and implementation*. Norwood, NJ: Ablex.

Kerr, S. T. (1986). Learning to use electronic text: An agenda for research on typography. *Information design journal*, *4*, 206–211.

Kinnear, P. R., & Gray, C. D. (2006). *SPSS 14 made simple*. East Sussex: Psychology Press.

Kol, S., & Schcolnik, M. (2000). Enhancing screen reading strategies. *Calico journal*, *18*, 10–14.

Laakso, S. A., Laakso, K. P., & Saura, A. J. (2000). Improved scroll bars. *Extended abstracts CHI 2000*, (pp. 97–98).

Landoni, M., Wilson, R., & Gibb, F. (2000). From the Visual Book to the WEB Book: The importance of good design. *Research and advanced technology for digital libraries*, (pp. 305–314).

Landoni, M., Wilson, R., & Gibb, F. (2001). Looking for guidelines for the production of electronic textbooks. *Online information review*, *25*, 181–195.

Lichtheim, M. (1973). *Ancient Egyptian literature: Volume I: The old and middle kingdoms*. Los Angeles: University of California Press.

Liesaputra, V. (2010). Finding information in a book. *Proceedings of the NZCSRSC'10*. Available at http://www.nzdl.org/Books/Books/nzcsrsc2010.pdf.

Liesaputra, V., & Witten, I. (2008). Seeking information in realistic books: a user study. *ACM/IEEE joint conference on digital libraries*, (pp. 29–38).

Liesaputra, V., & Witten, I. (2009). Computer graphics techniques for modelling page turning. *International journal on digital libraries*, *10*, 93–121.

Liu, Z. (2005). Reading behavior in the digital environment. *Journal of documentation*, *61*, 700–712.

Malama, C., Landoni, M., & Wilson, R. (2005). What readers want: A study of e-Fiction usability. *dLib Magazine 11*. Available at http://www.dlib.org/may05/wilson/

05wilson.html.

Marchionini, G., & Shneiderman, B. (1988). Finding facts vs. browsing knowledge in hypertext systems. *IEEE computer*, *21*, 70–80.

Marshall, C. C. (1997). Annotation: from paper books to the digital library. *Proceedings of the 2nd ACM international conference on digital libraries*, (pp. 131–140).

Marshall, C. C., & Bly, S. (2005). Turning the page on navigation. *JCDL '05: Proceedings of the 5th joint IEEE and ACM conference on digital libraries*, (pp. 225–234).

Marshall, C. C., Price, M. N., Golovchinsky, G., & Schilit, B. N. (1999). Introducing a digital library reading appliance into a reading group. *Proceedings of ACM digital libraries '99*, (pp. 77–84).

Marshall, C. C., & Ruotolo, C. (2002). Reading-in-the small: A study of reading on small form factor devices. *JCDL '02: Proceedings of the 2nd joint IEEE and ACM conference on digital libraries*, (pp. 56–64).

Martinez, A. (2010). Amazon.com's Kindle fails first college test. *The Seattle times*. Available at http://seattletimes.rwsource.com/html/businesstechnology/2011938870_kindle24.html.

McCusker, D. (1998). The design of Rocket eBook. Available at http://www.zippublishing.com/abstracts/rocket_ebook_ergonomics.htm.

McDonald, S., & Stevenson, R. J. (1998). Navigation in hyperspace: An evaluation of the effects of navigational tools and subject matter expertise on browsing and information retrieval in hypertext. *Interacting with computers*, *10*, 129–142.

McKnight, C., & Dearnley, J. (2003). Electronic book use in a public library. *Journal of librarianship and information science*, *35*, 235–242.

McKnight, C., Dillon, A., & Richardson, J. (1989). A comparison of linear and hypertext formats in information retrieval. *Proceedings of hypertext II*, .

McKnight, C., Dillon, A., & Richardson, J. (1991). *Hypertext in context*. Cambridge: Cambridge University Press.

McNamara, D. S., Kintsch, E., Songer, N. B., & Kintsch, W. (1996). Are good texts always better? Text coherence, background knowledge and levels of understanding in learning from text. *Cognition and instruction*, *14*, 1–43.

Mihalcea, R., & Csmoai, A. (2007). Wikify!: Linking documents to encyclopedic knowledge. *Proceedings of ACM conference on information and knowledge management*, (pp. 233–242).

Millburn Survey (2010). Flash player penetration. Available at http://www.adobe.com/products/player_census/flashplayer.

Mills, C. B., & Weldon, L. J. (1986). Reading text from computer screens. *ACM computing surveys*, *19*, 329–357.

Milne, D., & Witten, I. (2010). An open-source toolkit for mining Wikipedia. *Submitted to Artificial Intelligence*.

Morris, M. R., Brush, A. J. B., & Meyers, B. R. (2007). Reading revisited: Evaluating the usability of digital display surfaces for active reading tasks, . (pp. 79–86).

Murray, T. (2002). Metalinks: Authoring and affordances for conceptual and narrative flow in adaptive hyperbooks. *International journal of artificial intelligence in education*, *13*.

Muter, P. (1996). Interface design and optimization of reading of continuous text. *Cognitive aspects of electronic text processing*, *LVIII*, 161–180.

Muter, P., Latremouille, S. A., Treurniet, W. C., & Beam, P. (1982). Extended reading of continuous text on television screens. *Human factors*, *24*, 501–508.

Muter, P., & Maurutto, P. (1991). Reading and skimming from computer screens and books: The paperless office revisited? *Behaviour and information technology*, *10*, 257–266.

Nicholas, D., Huntington, P., Jamali, H. R., Rowlands, I., Dobrowolski, T., & Tenopir, C. (2008). Viewing and reading behaviour in a virtual environment. *Aslib proceedings: New information perspectives*, *60*, 185–198.

Niederhauser, D. S., Reynolds, R. E., Salmen, D. J., & Skolmoski, P. (2000). The influence of cognitive load on learning from hypertext. *Journal of educational computing research*, *23*, 237–255.

Nielsen, J. (1998a). 2D is better than 3D. http://www.useit.com/alertbox/981115.html.

Nielsen, J. (1998b). Electronic books: A bad idea. Available at http://www.useit.com/alertbox/9890726.html.

Nielsen, J., & Molich, R. (1990). Heuristic evaluation of user interfaces. *Proceedings of CHI'90 conference on human factors in computing systems*, (pp. 249–256).

O'Hara, K., & Sellen, A. (1997). A comparison of reading paper and on-line documents. *Proceedings of SIGCHI conference on human factors in computing systems*, (pp. 335–342).

Olsen Jr., D. R., Taufer, T., & A., F. J. (2004). ScreenCrayons: Annotating anything. *UIST '04: Proceedings of the 14th annual ACM symposium on user interface software and technology*, *6*, 165–174.

Parsons, C. G. (2001). *The efficiency and preference implications of scrolling versus paging when information seeking in long text passages*. PhD. thesis Department of Educational Technology, University of Northern Colorado, Greely, Colorado.

Poynter Institute (2000). EyeTracking online news. *Stanford-Poynter project*. Available at http://www.poynterextra.org/et/i.htm.

Quesenbery, W. (2001). What does usability mean: Looking beyond 'ease of use'. *Proceedings of the 48th annual conference of society for technical communication*. Available at http://www.wqusability.com/articles/more-than-ease-of-use.html.

Reynolds, L. D., Marshall, P. K., & Mynors, R. A. B. (1983). *Texts and transmissions: A survey of the Latin classics*. Oxford: Clarendon Press.

Ricker, T. (2011). Sony S1 and S2 dual-screen honeycomb tablets get official. *Engadget*. Available at http://www.engadget.com/2011/04/26/sony-to-launch-two-honeycomb-tablets/.

Rowland, F., Bell, I., & Falconer, C. (1997). Human and economic factors affecting the acceptance of electronic journals

by readers. *Canadian journal of communication*, *22*, 61–75.
Rowland, F., McKnight, C., & Meadows, J. (1995). *Project ELVYN: an experiment in electronic journal delivery*. London: Bower Saur.
Ryan, C., & Henselmeier, S. (2000). Usability testing at Macmillan. *Keywords*, *8*, 188–202.
Ryman, A. (2010). Profs: Kindle no threat to college textbooks. Available at `http://www.azcentral.com/business/articles/2010/07/06/20100706amazon-kindle-school-textbooks.html`.
Schcolnik, M. (2001). A study of reading with dedicated e-readers. Available at `http://www.planetebook.com/downloads/schcolnik.pdf`.
Schilit, B. N., Golovchinsky, G., & Price, M. N. (1998). Beyond paper: Supporting active reading with free form digital ink annotations. *Proceedings of CHI'98*, (pp. 249–256).
Schilit, B. N., Price, M. N., Golovchinsky, G., Tanaka, K., & Marshall, C. C. (1999). As we may read: The reading appliance revolution. *IEEE computer*, *32*, 65–73.
Sebrechts, M. M., Vasilakis, J., & Miller, M. S. (1999). Visualization of search results: A comparative evaluation of text, 2D and 3D interfaces. *Proceedings of the* 22nd *SIGIR conference on Research and development in information retrieval*, (pp. 3–10).
Severinson, E. K., Fatton, A., & Romberger, S. (1996). The paper model for computer-based writing. *Cognitive aspects of electronic text processing*, *58*, 137–160.
Shneiderman, B. (1989). Reflections on authoring, editing and managing hypertext. *The society of text: Hypertext, hypermedia and the social construction of information*, (pp. 115–131).
Shneiderman, B. (1998). *Designing the user interface: Strategies for effective human-computer interactions*. Menlo Park, CA: Addison-Wesley.
Spool, J., Scanlon, T., Schroeder, W., Snyder, C., & DeAngelo, T. (1997). *Web site usability: A designer's guide*. North Andover: User Interface Engineering.
Suetonius (2000). *Lives of the Caesars*. USA: Oxford University Press.
Swan, R. C., & Allan, J. (1998). Aspect windows, 3D visualizations, and indirect comparisons of information retrieval systems. *Proceedings of the* 21st *SIGIR conference on Research and development in information retrieval*, (pp. 173–181).
Tashman, C. S., & Edwards, K. W. (2011). LiquidText: a flexible, multitouch environment to support active reading. *CHI 2011*, .
Tavanti, M., & Lind, M. (2001). 2D vs. 3D implications on spatial memory. *InfoVis 2001 Symposium on Information Visualization*, (p. 139).
Waycott, J., & Kukulska-Hulme, A. (2003). Students' experiences with PDAs for reading course material. *Personal and ubiquitous computing*, *7*, 30–43.
Williams, G. (2000). A few things that absolutely must change in the e-Book industry. Available at `http://www.ebooknet.com/story.jsp?id=2232`.
Wilson, R., Landoni, M., & Gibb, F. (2003). The WEB Book experiments in electronic textbook design. *Journal of documentation*, *59*, 454–477.
Wingo, E. O. (1972). Latin punctuation in the classical age. *Janua linguarum seris practica*, *133*, 164–166.
Zhou, B., Luo, P., Xiong, Y., & Liu, W. (2009). Wikipedia-graph based key concept extraction towards news analysis. *2009 IEEE conference on commerce and enterprise computing*, (pp. 121–128).

Appendix: Questions used in the evaluation tasks

Browsing task (Section 8.3). Some questions are followed by a list of possibilities that depends on the book itself.

1. What kind of book is it?
2. Which of the following would make the best title for the book?
3. Who do you think this book was written for?
4. For each of these topics, write down the section numbers that explain it.
5. Put a check next to each topic discussed in the book.
6. Put a check next to each statement that reflects the underlying beliefs or point of view of the book.

Searching task (Section 8.4). Includes whether the question should be answered from the subject index or by search, and whether it is factual or analytical.

1. Who created the CART algorithm? (*subject-index, factual*)
2. Write down the page numbers that explain a technique used to analyze people's buying patterns in a grocery store. (*subject-index, analytical*)
3. In statistics, there are four types of attribute values. Choose two that are used by the data mining system. (*text-search, analytical*)
4. Write down the page numbers that explain a situation where prior knowledge can be used to cut down a rule set. (*text-search, factual*)

Enriched document representation task (Section 8.5). All questions are followed by a list of possibilities from the book.

1. Match the following terms with the specified list of definitions.
2. For each of these topics, down write the page numbers that explain it.
3. Put a check next to each topic discussed in the book.

Personalization task (Section 8.6).

1. Mark at least two passages in each subsection of the chapter.
2. For each marked passages put, in a note, the reason for marking it.
3. Bookmark every page that contains a figure.
4. At the end of each chapter, add a note containing at least a sentence or three keywords that best describes it.