

Who was that? Enriched eReading through in-book visualisation

Annika Hinze, Claire Timpany, Nicholas Vanderschantz, Stuart Thomson
Computer Science Department, University of Waikato, New Zealand
{hinze, ctimpany, vtwoz}@waikato.ac.nz, swdt1@students.waikato.ac.nz

The shift to electronic books has opened a realm of possibilities for assisting users in reading books, but little of this potential has been realised. Today's eReaders offer simple search features but rarely make any ongoing supportive data about a book's content available. We investigated the use of visual aids for eBook reading and developed an eReader prototype. Our prototype provides a timeline of character appearances throughout a book, coupled with interactive identifications and lists of characters as they appear within the narrative. We present results of a user study in which we observe that contextually-relevant information can enhance a reader's experience.

eReader, eBook, Digital library, Semantic analysis, BookLine, CharacterChart, Characters, Narration

1. INTRODUCTION

This paper explores the use of visual aids for eBook reading. If you ever returned to continue reading a book after a long break you are probably familiar with the question “who was that?” when forgotten character re-appears. You may wonder when they last appeared in the story or who they appeared with. We developed an eReader application that provides information about the characters as they appear in works of fiction. The goal is to help readers keep track of the character ensemble and their appearances throughout the book, using data extracted from the book itself.

Our enhanced eReader visualises how often a character appears at various stages of the book using a line chart that we call the *BookLine*. We also provide interactive in-page features that notify readers of the appearance of characters as they appear within the narrative, which we call the *CharacterChart*. This CharacterChart works hand-in-hand with the BookLine. This paper describes the design and implementation process followed to develop our enhanced eReader prototype and the user evaluations that we have conducted to date.

We previously explored a number of options for visual aids in eReaders (Coleman & Hinze, 2017) and found that any information or visualisation that obstructs the text, e.g., as a pop-up, is a hindrance to the reading flow. The work reported in this paper builds on our previous research in the following way. To extend the concepts developed earlier, we observed that traditionally physical books on paper or parchment were often annotated in the margins, to help the reader track the content. We have used

this feature as a design inspiration for our creation of a visually annotated prototype eReader.

A range of design considerations were taken into account, such as being conscious to not spoil the plot of a book by early reveals of characters, or by overly summarising a narrative. The interface and interactions were designed such that they did not disrupt the flow of reading, yet still provided relevant contextual information and were suitable for mobile reading on small screens.

We evaluated the CharacterChart and BookLine features in a user study in which participants engaged with the eReader for a period of time. We present here the results of this initial study and discuss the reception of our eReader and possible extensions and adaptations for future research. The remainder of the paper is structured as follows: Section 2 provides an overview of related work and design inspirations. Section 3 introduces the design of our eReader prototype. Section 4 outlines the evaluation and Section 5 summarises the results of our user study. In Section 6, we discuss the insights from the study results and possible future research adaptations of our concepts. The paper concludes in Section 7 with a short summary.

2. RELATED WORK

In this section, we discuss related work that is concerned with ways of presenting auxiliary information in both physical books and eBooks.

Auxiliary information in printed texts. The Jewish Talmud is a collection of writings that were compiled

over centuries, combining original texts and commentaries in the margins around the text. As such the Talmud illustrated a solution to an early example of the problem faced by print media, where extra information is challenging to display alongside a text. Over time, as more commentaries were added and space in the margins ran out, the original text was made smaller, so that more commentaries could be added. Today, typically more space is taken up by the commentaries on a Talmud page than the original text (Fishman, 2011). Our work is inspired by this use of extended margins for displaying auxiliary information.

Notes in physical books. A common solution for showing more information in books is the use of footnotes, endnotes and sidenotes. This information is often presented at the bottom of the page (footnotes) or at the end of a chapter or book (endnotes); sometimes the margins are used as well (sidenotes). An indicator in the main text usually refers the reader to the note. Some nonfiction texts incorporate notes into an additional column. This allows readers to access additional information while continuing to read, or to find extra information in a predictable location (Mitchell & Wightman, 2005). Sidenotes and additional columns are sometimes used in a range of non-textual ways. For example, Tufte (2006) recommends incorporating small visualisations embedded in the margins which may avoid having to 'jump' between text and annotation.

Cross-references in printed texts. Timpany (2012) explored the use of cross references in physical books. Her designs were inspired by concepts of hypertext (i.e. navigating between physical pages with physical 'hyperlinks'). She created a series of physical book prototypes, in some of which marginal annotations encoded cross-references. Other prototypes required the reader to interact with the book to reveal the hidden annotations, such as the use of transparent sheets with annotations overlaid on the page (annotations shown in the previously-empty margin space).

Auxiliary information in web documents. Information on websites is often presented dynamically, displaying additional information when a reader activates an element. A common technique is a hovering event, in which new information appears when a user places their mouse cursor over an element for a short time. The additional information is often presented as a pop-up element on top of the existing content.

The knowledge graph (Singhal, 2012) is a feature of the Google search engine in which a summary is provided of relevant information to a search. The aim of this feature is somewhat similar to that of our project presented here, as we also aim to present readers with relevant information about a digital (book) text. The knowledge graph appears to the

right of the search result list, in a space that is typically less full (but not necessarily empty). As such it appears in a space that is visually similar to extended margins.

Enrichment of eBooks for children. Visual enhancements through illustration are known to support better understanding of a text (Sipe 1998; Fang 1996). Colombo et al. (2014) proposed six design guidelines for more engaging children's eBooks. One of these recommends to make books easier to understand through the use of inline dictionaries, illustrations, and additional text. Another guideline recommends the enhancement of the text through visual and illustrative techniques that reinforce concepts in the text. It is typically stressed that enhancements should never replace the text but enrich and illustrate the text's meaning.

Semantic and visual enrichment books. Coleman & Hinze (2017) developed an eReader to investigate the types of additional book information that may be presented to the readers of an eBook. Their goal was to provide contextually relevant information that would be easy to access. They focussed on works of fiction and particularly on providing auxiliary information about relationships between characters. Their application performed automatic semantic analysis of the text of an uploaded eBook. Words (often characters) for which further information was available (either from their analysis or from identified external sources) were then highlighted and made interactive. On selecting one of the interactive words, a reader could then choose from a variety of information to be displayed, such as relevant Wikipedia pages, character occurrence, proximity to other characters or terms throughout the book, and clusters of verbs or nouns associated with a selected word through part-of-speech analysis). They observed that when too many terms were highlighted, this potentially distracted readers. Readers were concerned that presented information may reveal unwanted information about the story (e.g., if a character disappears prematurely or re-appears at the very end). Some of the semantic information was described as being useful for someone who is interested in analysing the book in more detail but may be unwanted for a new reader. Their study results revealed that a number of participants used the provided information to avoid needing to read the whole book. This went counter to their goal of complementing a book and supporting a reader instead of replacing the reading.

Information enrichment in commercial systems. A number of commercial eReaders incorporate features that extend the traditional reading experience. Text within an eBook on a Kindle Paperwhite is selectable, then allowing access to dictionary, translation, and Wikipedia search services. Web based and hardware independent eBook readers, such as Google Play Books and

Apple iBooks, implement footnotes as floating containers overlapping the text. These systems require interaction from the reader to access the additional information. While information in these systems is kept off-screen until selected, it often is implemented in such a way that it overlaps and obscures the book's main content when selected by the reader. Similarly, the Aldiko eReader provides tools for highlighting text, creating and sharing notes, performing dictionary lookups, and text searches of the book. Various other eReading systems also offer tools such as text-to-speech, web search, as well as traditional orientation and navigation features similar to physical books.

To the best of our knowledge, there are no other systems that aim to provide semantic information related to a book's characters and story line beyond the research performed by Coleman & Hinze (2017) and the work described in our paper.

3. DESIGN & IMPLEMENTATION

This section discusses the considerations during designing and developing our eReader software and its interface. We start by illustrating our prototype eReader and the concepts of the BookLine and the CharacterChart. In order to explore eReader features, we developed our own eReader software package. For easy access via web and mobile devices, it was implemented as web service. Our eReader implementation is available as open source software package on <https://github.com/s-thom/520-reader> and an online software prototype is available on <https://reader.sthom.kiwi>.

For developing, testing and evaluation of our eReader application, we used the book *Through the Looking Glass* by Lewis Carrol (Carrol, 1871), which is made available via Project Gutenberg.

3.1 Character Analysis

We used Named Entity Recognition to find references to book characters within the text. When the application loads, first a list of all the characters is retrieved. Each character in the list has a display name, a list of alternate names, and optionally an image representing that character. As an example of alternate names, Alice's cat Dinah is called "Dinah" and "old cat" in *Through the Looking Glass*. When a book page is displayed in the eReader application, the software iterates through the list of characters. If a character's name or alternate name is found on the current page, that character is added to a subset of known characters up to this point in the book. This list known characters is then used for creating both the BookLine and CharacterChart. A screenshot of the eReader application is shown in Figure 1.

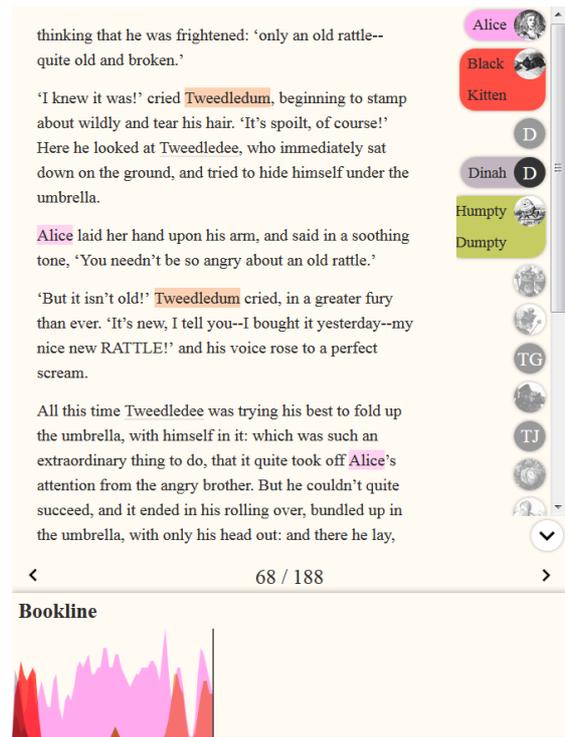


Figure 1: Screenshot of eReader

3.2 CharacterChart

Throughout the book, each character's name is given a grey underline within the text displayed in a page. This underline aims to assist the reader in locating characters in the page when they skim-read the text, and is also intended to reinforce the presence of a character during reading. We couple this in-page identification of characters with a chart of known characters that is displayed on the right-hand side of the application, called CharacterChart (see Figure 1 and 2). Characters appear in this chart in alphabetical order, to maintain consistency between pages. The CharacterChart displays only characters who are mentioned on the currently opened page or on any previous pages in the book.

The CharacterChart is an interactive page element that can be used visually to remind readers of the characters present on the page at a glance. Characters on the CharacterChart can also be tapped by the user to select that character. For a selected character, their CharacterChart entry and all in-page instances of that character are highlighted. Highlighting changes the grey underlining to a full-colour background of each instance of their name in the text and a highlight appears behind the character's icon in the CharacterChart. For each character, we use the same highlight colour for the names on in the page as well as on the BookLine and in the CharacterChart. This highlighting creates a visual connection between these three aspects of the interface. In the example shown in Figure 2, Alice

and Tweedledum are highlighted and Tweedledee is present but not highlighted.

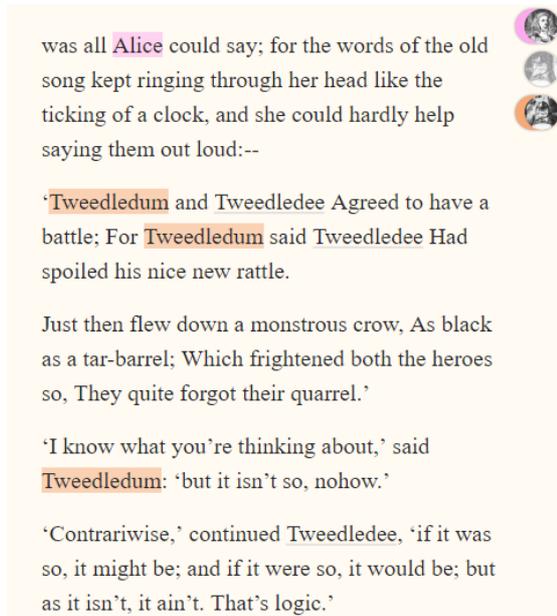


Figure 2: Coloured highlighting of selected characters, unselected characters are underlined.

3.3 BookLine

The use of a line graph to display the frequency of appearance of a character in parts of a book was previously investigated by Coleman & Hinze (2017). Similar to their character frequency graph, our BookLine device is intended to provide a visual representation of the frequency with which the author references an individual character throughout the book. While they had tested a number of visual elements, we here explore the BookLine feature in much more detail.

3.3.1 BookLine concept

The BookLine is an interactive page element that can provide a visual representation of the characters in the story read so far. It also serves as a navigational feature, so a reader can easily move to sections of importance for a particular character within the book. This navigation is achieved by tapping any point on the BookLine.

We use a continuous line graph in which high point indicates a character having been referenced often in a section, while a low point on the graph represents a section where the character is a minor player in the narrative. During development we observed that the bookline turned out often very rough with many sharp points. While this would show an accurate representation of the character's appearances throughout the book, it can be difficult to identify an overall trend. In order to make the BookLine easier to interpret, we smooth the values using a weighted moving average. This takes the pages surrounding the current page into account, but giving more weight to the pages closest to the

current page. The result, shown in Figure 3, is a much smoother line. This makes it easier to see the overall trends in character appearances. It also removes the need to filter out noise that is created through pages of varying length.

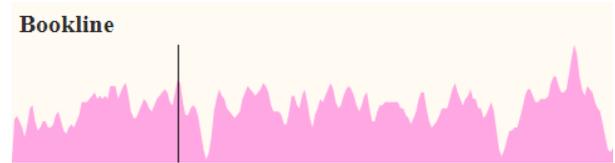


Figure 3: The smoothed BookLine.

3.3.2 Avoiding Spoilers in the BookLine

An important consideration when displaying character information in the BookLine is that the readers should not encounter any information before it appears in the text. In developing the BookLine visualisation, it needed therefore to be carefully considered how the BookLine is made available to readers, to avoid plot spoilers. For example, if we were reading a murder mystery book and knew that a character stops being mentioned after a certain point, it may be highly likely that the character is murdered at that point. This could ruin, or spoil, the surprise and development in the plot.

While character frequency information, as shown in the BookLine, does not explicitly give any information about the events happening in the book, it may show when characters directly interact, or at least appear in the book in close succession. In order to prevent these plot spoilers of future character appearances, we have limited the BookLine to only show character frequency for pages that have already been read (see Figure 4 and Figure 5).



Figure 4: BookLine only showing to the current page.



Figure 5: The BookLine displays pages that have been read, even when the reader navigates backwards.

This prevents the reader from being able to see what is yet to happen in a book. In addition, this implementation of the BookLine provides the reader with an indication of how far through the book they have read. This is an unintended affordance of our eReader providing reading orientation cues that are

often missing from eReaders and eBooks presently (Vanderschantz et al., 2017).

3.3.3 Spacing considerations for the BookLine

The visualisation of character frequency implemented by Coleman & Hinze (2017) used information overlays. By overlaying the information on the page text, the reader was unable to continue reading the text while viewing the information. While their eReader focused on displaying on-demand semantic information to the reader, our application aims to assist the reader without (much) interrupting their reading flow.

When designing the prototype, we wished to ensure that readers can continue to engage with the book without being interrupted. In early prototypes, the BookLine was brought up from the bottom of the screen without changing any other part of the display. However, this meant that the bottom portion of the page was covered and the reader could not continue to read the text and view the BookLine at the same time.

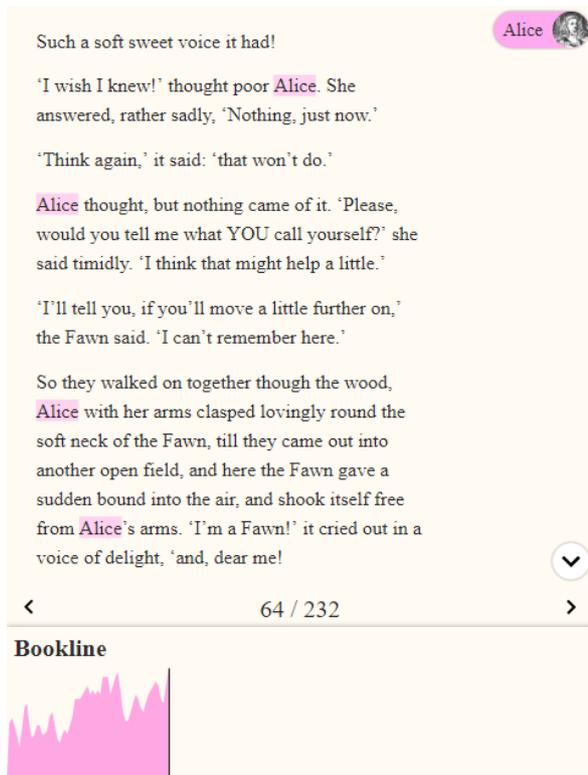


Figure 6: The BookLine for Alice, with text area scaled.

Our solution was to shrink the text area to allow space at the bottom of the eReader page at the same time as the BookLine comes into view. This also made space available at the right side for the CharacterChart (see Figure 6). This means the reader can continue to engage with the text while referring to the BookLine for information and when the BookLine is not present, there is simply a larger bottom margin present in the eReader.

3.3.4 Multiple Characters on the BookLine

Important to the function and usefulness of the BookLine was how to display multiple characters' BookLines at the same time. This allows the reader to compare how often and where in the narrative one character appears in contrast to another character.

Showing multiple BookLines is intended to also give the reader a sense of the relationships between characters. If two characters are mentioned near each other in the text, then their BookLines will be similar or at least overlapping (see for example Figure 7).



Figure 7: BookLines for Tweedledum and Tweedledee

If there is no relationship between the characters, then their instances on the BookLine would not occur in the same places, meaning their individual BookLines do not or barely overlap (see Figure 8).



Figure 8: BookLines for Tweedledee & Humpty Dumpty

We hypothesise that comparing characters in the BookLine could allow the reader to draw conclusions about a characters' interactions based on the readers prior knowledge of the narrative. Therefore, this visualisation may assist with narrative recall without the need to explicitly remind the reader what happened earlier in the book. This would also mean that the BookLine may offer information value without the reader determining the outcomes of important events without reading the book.

3.4 Using Colour

Our eReader application used colour to reinforce the link between the characters in various parts of the application. The use of colour is an important factor in the design of information graphics (Katz, 2012). At any given time, a character may be in the text, and shown in the CharacterChart on the side, be represented in the BookLine at the bottom of the page, or any combination of the three. Colour is a natural way of linking these distinct user interface elements together and distinguish between different characters in the BookLine.

Should it be necessary to algorithmically automate colour allocation, the colour of a character can be selected from a list of predetermined colours. In

initial development of this algorithm, we assigned the next “unoccupied” colour when a new character is selected. When deselected, that colour was made available again. However, during testing we realised that this method may lead to confusion, as characters may appear in different colours depending on the order in which they are selected. This resulted in a lack of learnability for the reader.

Instead, the current implementation allocates a different (but fixed) colour to represent each character in a book. Because of the large number of characters in the book we used, the colours of some characters appear visually similar, especially when shown on the BookLine. While this can cause some ambiguity between characters’ appearances, it also allowed us to represent the closeness of character relationships by colour as well. In *Through the Looking-Glass*, the Red Queen and Red King are on the same side of the game of chess. Because of this, and the colour in their names, we gave them both a shade of red, varying the brightness so they could still be distinguished. This meant that more visually distinct colours could be used for other characters, or groups of characters. The colour for each character is included in the character information that is loaded when the application starts.

This problem of colour assignation is an interesting design problem. Colours maybe preferable to be assigned such that characters who appear often together have colours that are very distinct to ensure clear readability of the highlights. Furthermore, indicating closeness of characters in colour may provide further spoilers as to the future development of a story. Whichever way, for commercial implementations this could mean that when budgets allow, a designer would implement colour decisions for book characters to ensure strong semantic relationships while maintaining visibility. In other circumstances the algorithm would assign colours.

4. STUDY METHOD

To evaluate the success of our enhanced eReader, we performed a user study with 13 participants to investigate if the visualisations employed in our system assist in deeply reading and exploring books. Participants were supplied with our eReader software and the freely available book *Through the Looking-Glass* (Carroll, L. 1871).

The eReader was accessible by the participants via a web browser on any device they chose to use. Participants were encouraged to read the eBook in their own time. The participants’ direct interactions with the application were recorded by our eReader system: including page turns, character selection page number visits, and time stamped recording.

Interviews. Once they had completed reading the book the participants were invited to take part in

short semi-structured interview to elicit their feedback about the application. In this interview participants were asked to name three BookLine aspects they enjoyed or found useful as well as three suggestions for improvements. The questions were open ended to elicit a wide range of answers. During the interview, extensive field notes were taken.

Participants. 13 participants were recruited through personal referral and self-nomination. Five of the 13 participants were female (P3, P4, P8, P9, P11) and 8 were male (P1, P2, P5, P6, P7, P10, P12, P13). 10 participants were students. Most participants had never read the book *Through the Looking-Glass* before we asked them to participate.

5 RESULTS

We present here the results of our users’ interaction with our eReader software followed by our interviews with all 13 participants. Due to a technical issue, only data for eight participants was reliably recorded by the test apparatus. Therefore, the data reported in Figures 9 through 11 reports only the results of nine participants. We report the results of interviews with all 13 participants.

Figure 9 shows how far through the book each of the nine participants progressed. Only participants P7 and P9 read less than 50% of the book.

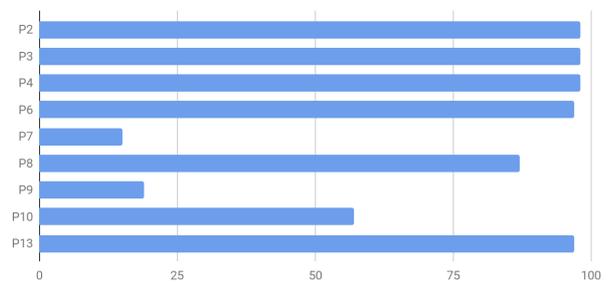


Figure 9: Reading progress (% of pages)

Figure 10 shows how often each character was selected by participants to be highlighted. This included participants tapping a character in the CharacterChart and tapping a character in the text.

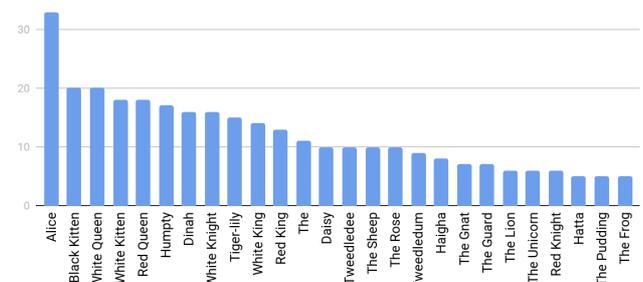


Figure 10: Selection of character

This shows that Alice, the main character in the novel was selected more than 30 times by the eight participants who we were able to capture data for.

Figure 11 shows how often other characters were selected by the eight participants to be shown together (or in contrast to) the story's main character, Alice.

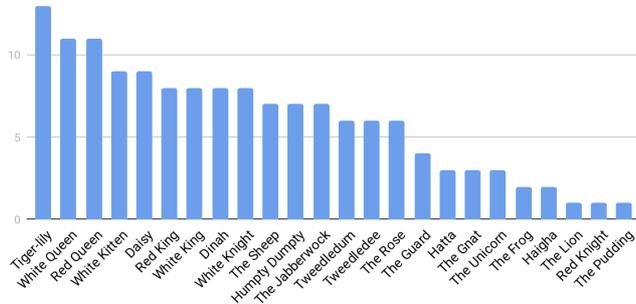


Figure 11: Comparison to Alice

Figure 12 shows aggregated access to the BookLine feature mapped across length of the book. The most frequent use can be seen at the start of the book (70 accesses for 8 participants), presumably when users were becoming familiar with the use of this feature. An additional spike at about 10-20% into the book is where Alice talks to the Red Queen in the flower garden.



Figure 12: BookLine access across the book

The book introduces some new characters here, and also has its first character reappearance of the Red Queen. The next area of activity is at ~35% where Alice meets Tweedledum and Tweedledee. Use of the BookLine continues as characters are introduced, until the very end where there we observed a higher number of character interactions. We observe that participants continued to use BookLine features throughout their reading and the correlation between character appearances and interactions in the story appeared to induce higher use of the BookLine.

5.1 User Experience

Here we outline the feedback participants gave about their experience with the eReader.

Several participants mentioned how the BookLine helped them explore parts of the book they had already read. One participant recalled how Alice and the Red Queen met towards the end of the book, and in using the BookLine this participant rediscovered the earlier interactions between the two characters. Using the BookLine allowed this participant to better understand the relationship between the Red Queen and Alice.

Two participants stated that the BookLine would be useful in other books they were currently reading. Both participants noted that they sometimes forget important detail about the characters between reading sessions or as the narrative progresses. These two participants described how if they were using the BookLine, they would be able to see where the characters appeared previously in the book, and use that to remember details about the characters in those other books.

5.2 eReader enhancements

Here we discuss the eReader enhancements and features that attracted feedback from participants.

5.2.1 The CharacterChart

P1, P7, and P8 explicitly discussed how useful the CharacterChart was for their reading. They found the list made it easy to recognise which characters they had already seen in the book. The colour coding of the characters helped the participants distinguish between characters. P1 explicitly stated that the “icons on the side are useful”, while P2 reported the “colour coding is useful” in disambiguating which character is which. However, P1, P6, P7 found the CharacterChart difficult to use. They felt the icons were too small and they found it difficult to distinguish which image belonged to which character (without having to bring up the name). P6 suggested the CharacterChart was “too distracting” because it was “difficult to know which character was which [without] hover or mouse over labels”.

Participants recommended methods to enhance and improve the CharacterChart. P1 suggested a search functionality to find the character in the list, while another participant noted the absence of a hover text. P8 suggested mixing the image with the character's colour, to reinforce the connection between the colour and character. P8 also noted that not all characters had associated images and the use of (the same) initials caused confusion. Hover and tooltips were also noted as ways to improve both the CharacterChart and the BookLine.

5.2.2 Highlighted characters

P3, P5, and P10 commented positively on the highlighting and underlining of character names in-page text. In particular participants discussed the ease of identifying characters in the text when they

were underlined or highlighted. P6 relayed a situation in which they wished to find the name of one of Alice's kittens and used the underlining. P6 said "I liked how nicknames worked for Snowflake the white kitten". P6 felt without the underlining they would have missed this detail.

P5 described their dyslexia as making reading continuous text a considerable effort. They reported using the CharacterChart to highlight characters on each page, and experimenting with both underlined character names and highlighted character names in-page text. P5 suggested that by using these text enhancements they were able to find the start of the next line to read with greater ease. P5 stated "it helped me find my place in the text". Further, P5 reported that the underlines and highlights reduced their reading time of individual pages by over 50%.

5.2.3 The BookLine

Eight participants commented on finding the BookLine interesting or useful and particularly mentioned the character frequency information. One participant (P4) described how the BookLine chart helped them to understand what had happened previously in the book, stating that when they were revisiting the book the BookLine served as a reminder of "oh, that's where that happened". P3 similarly stated that the BookLine was "useful if you forgot characters" either when a character re-appeared later in the book or when they had been interrupted in reading the book for a while.

Both P5 and P9 found the BookLine enjoyable to "play with". However, these two participants wished to engage with this feature so much that it became a distraction from reading. P6 decided to stop using the BookLine entirely so they could continue read the book without distraction.

P8 reported that the BookLine was a good size stating that the "BookLine was broad enough across the screen to show dips and troughs" in a character's line, while P9 found it to be too small.

P8 found the BookLine useful when comparing the frequency with which characters appeared throughout the story. Though when many characters were selected the BookLine became confusing for P8, who stated that "any more than five characters [on the BookLine] is confusing". P9 attempted to use the eBook on a mobile device and found the BookLine too small to use on this small screen.

P12 struggled to understand the use for the BookLine when first presented with the eBook. They stated that "I had to figure out that the line was character frequency". This participant further noted that they had received too little instruction on how to use the eBook and could not see the relevance of the BookLine for their needs as a reader. Similarly, P11 also suggested building a tutorial into the eBook or app to assist with use of these enhancements.

P11 stated they did not use the BookLine as it was not shown to them and they did not discover it during their time using the eBook for the study.

Another participant reported that while they appreciated seeing where the characters appeared in the book, they were missing sufficient context in the BookLine chart. They wished to know how many character mentions were represented by a spike in the BookLine chart.

5.3 Interface and Interaction

Participants also commented on aspects of the user interface and user interaction.

5.3.1 Colours

Participants commented that many of the colours used for the various characters were too similar (e.g., Red King and Queen), potentially leading to confusion. One participant commented that they almost completely ignored the colours in the CharacterChart, instead only reading the names of the characters in the "bubble". P7 stated that "some colours for characters are very similar" recommending some further work in refining this differentiation for users.

5.3.2 Device Sizes

Two participants commented on the size of the BookLine on their devices. The first participant used a tablet device to read the book, and found that the BookLine was a good size for them to use. The second participant used their smaller smartphone to read the book, which then reduced the size of the BookLine to a point where it was more difficult to use. Both devices had a similar aspect ratio, which meant that the size of the BookLine relative to the viewport size was the same. However, due to the large difference in physical size of the device, the actual size of the BookLine was greatly reduced on the smartphone.

5.3.3 Further interface and interaction feedback

Some of the usual software hygiene issues were discovered during the study. Issues discussed included insufficient reflow when a user switched between full-screen and a smaller browser window. Another user discovered limitations caused when the margin size was reduced due to browser window sizing. This smaller margin resulted in shortening of character names in the CharacterChart that caused character identification issues for this participant (missed recognitions of shortened names, such as "unicorn" and "queen"). One participant tried using the application on two differently-sized screens and discovered continuity issues with the last-visited-page (the app opened on the correct page number but different text was presented). These typical issues were outside the scope of the study being undertaken yet show the depth with which the

participants considered and used the enhanced eReader during this study.

One participant suggested it would be useful to have the option to completely disable all interactive features, allowing the user to read the book with no potential for distraction.

Three participants wished for additional online help in using the system, while many others found the application intuitive to use.

6. DISCUSSION

The results of the user study suggest that visualisation of the narrative could enrich the reading experience during eReading. As well as using individual character timelines, participants in our study used the BookLine to compare characters and reported better understanding of the book because of it. Additionally, the CharacterChart and in text highlighting and underling was discussed positively by participants. Here we discuss selected aspects of the participant feedback and their relationship to the work of others.

6.1 Semantic enrichment

Previous work on semantic enrichment found that too much aggregation of information created issues with potential plot reveals and reduced engagement with the book. This was also shown in the user study conducted by Coleman & Hinze (2017) where participants discussed using the visualisations to get the required information rather than reading of a text in its entirety. Given this, the methods of visualisation that we have set out to develop for this project were intended to enhance the reading experience rather than replace it. No participant reported plot spoilers as a feature of their reading experience suggesting that we have been able to ensure this through our implementation of staged revealing of the Bookline.

6.2 Play and distraction

Many participants reflected on using the enriched eReader features in terms of both play and distraction. When developing our eReader we intended the enhancements to serve informational purposes and provide deeper understanding of the narrative for readers. Interestingly, participants used the BookLine for fun and exploration and not solely for informational purposes. Comparing character timelines gave participants the sensation of discovering new knowledge. These results support the findings of Sirkia and Sorva (2015) providing further evidence for the need for research into visual representations of information in eReading.

Conversely, one participants also noted that they did not use the BookLine because they felt it was

distracting from their reading of the story. Aspects of enjoyment and playful interaction compared to distraction warrant future study.

6.3 Addressing observed issues

Increasing the height of the BookLine would reduce the available space for the text area, making it more difficult to read while the BookLine is visible. In addition, the BookLine would not be as easy to interpret as the points would appear more extreme. A different solution would be to have the application only show a portion of the BookLine at a time. However, this means the user does not have the context of the entire book, such as how far through the book they are.

Shortened names (e.g. "Queen" vs "Red Queen") and pronouns referring characters ("Alice" vs "she") are not recognised as a character occurrence in the current software version. This could be remedied by using a semantic-enhanced search or PoS tagging, as described in Coleman & Hinze (2017).

We had purposefully not included additional numbers or a legend in the BookLine chart as the lines are presented in a smoothed fashion and our focus was on overall impression vs statistics tool. We fear that adding such detailed information would detract from the supportive function of the chart (vs a statistical analysis of the text) and would present a false impression of accuracy. The frequency of mentioning of a character's name is only indicative for the importance of a role (especially as only their full name is recognised). Furthermore, an author's style in using names or pronouns strongly influences these statistics and a comparison between books could seem desirable (if numbers were presented) but would potentially give a false impression.

While the ability to use colour to help distinguish characters or indicate relationships between characters was considered in the design stage (see section 3.6), several participants commented that the closeness of colours was problematic for them. Research suggests that we are only able to clearly distinguish 5 to 7 colours depending on context and saturation or brightness (Katz, 2012; Shneiderman & Plaisant, 2004). The colours used, in an interface, such as the Bookline, should be clearly distinguishable from one another and be easily remembered and associated (Katz, 2012). The strong reliance on colour here may also be an issue for colour-blind readers wanting to use the BookLine. Future solutions for choosing methods to connect the characters between the CharacterChart, highlighted text and BookLine should also take colour-blind users into consideration.

The majority of recent research relating to colour in interface design is related to preference and aesthetics, (cites), our study therefore reinforces the

need for further investigation into the effective use of colour in interface design, especially when it forms a key part of the usability of the interface. Vanderschantz et al., (2010) identify that readers have both emotional and physical responses to colour and the use of colour can be motivational both in print and on screen.

Additional information is being presented in physical books, eBooks and online pages. Many digital media, such as websites and eBooks, overlay the information over the content. Physical media predominantly uses the space around the content for supplementary information, potentially reducing the space for the main content. Some online systems also use these increased margins techniques to present information in a way that does not obscure existing text.

6.4 Readers as learners and readers with learning difficulties

Dyslexia provides a number of difficulties for readers and learners. One of which may be difficulty finding the next line of the text when reading continuous prose. The unanticipated discussion with one participant who self-identified as having dyslexia suggests that visual features implemented in our enhanced eReader assisted him in his reading. The features that assisted this participant were the visual enhancement of underlined and highlighted character names in the page text. The participant believed these text enhancements helped him to identify the position in the text that he was reading. Significant investigation across many disciplines (see for example Parilová et al. 2016; Jing et al. 2017) is being undertaken to survey the potential for assisting readers and learners with various difficulties. Our focus here is presently not to investigate learning difficulties as an early outcome of this research, however we welcome the possibility that informational enrichment might assist readers with learning difficulties.

Our work here did not specifically investigate eReading enrichment for learners. However, it is widely reported (i.e. Sipe 1998; Fang 1996) that illustrative enrichment in children's books support children learning to read. With participants in our study describing use of the BookLine to reinforce understanding of the relationship of characters and storylines in the narrative investigation of how this type of enhancement could serve young readers during their education and literary students in high school and university education is warranted.

7. CONCLUSIONS

We investigated the use of eReading enrichments through enhancements to an eBook that provides contextually relevant information to readers.

7.1 Summary

In this paper, we introduced and evaluated the concept and implementations of the BookLine and the CharacterChart. The BookLine is an intuitive visualisation method for character frequency information, and its presence in the application increased the readers' enjoyment of and engagement with the book. The CharacterChart gives quick access to character information on the current page.

Participants of our user study reported that the additional interactive features (such as the CharacterChart on the side and the highlighted character names and the BookLine), assisted them and improved their reading experience. This initial user study shows that it is possible to enhance eReading to reinforce the connections present in the narrative and to create a cohesive experience for the reader appears possible. Even when readers described that they were merely "playing" with the BookLine, these readers were still engaged with the book. This playing resulted in readers exploring the relationships between the characters to discover information they had not gained through the act of initial reading. One participant reported that the eReader helped them read the book even though they are dyslexic. This aspect of the highlighting will have to be explored further in future work.

7.2 Future Work

One area for future work is integration of further methods for extracting character information. A semantically-enhanced character search could detect character references that are missed in the current application.

The work presented here purposefully focussed on displaying character information through the BookLine coupled with the CharacterChart, and in-page text highlighting. Further explorations are planned for other visualisation methods, together with a comparison of their effectiveness at providing useful character and narrative information while ensuring readers remain engaged with the book.

Another consideration is our use of colours for distinguishing characters. We wish to ensure the application is still usable by those with colour blindness. We are also interested in further exploring the usefulness of BookLine and other book visualisations for people with dyslexia and other learning difficulties.

8. REFERENCES

- Carroll, L. (1871). *Through the Looking-Glass* London, England: Macmillan Publishers Ltd
- Coleman, M., & Hinze, A. (2017). Visual semantic enrichment for eReading. In *Proc 31st British*

- Human Computer Interaction Conference* (pp. 6 pages). Sunderland, UK.
- Colombo, L., Landoni, M., and Rubegni, E. (2014). Design guidelines for more engaging electronic books: Insights from a cooperative inquiry study. In *Conf. on Interaction Design and Children (IDC)*, pages 281–284. ACM.
- Fang, Z. (1996). Illustrations, text, and the child reader: What are pictures in children's storybooks for?. *Reading Horizons*, 37(2), 3.
- Fishman, Talya. *Becoming the People of the Talmud Oral Torah as Written Tradition in Medieval Jewish Cultures*. 1st ed., University of Pennsylvania Press, 2011.
- Jing, C. T., & Chen, C. J. (2017). A Research Review: How Technology Helps to Improve the Learning Process of Learners with Dyslexia. *Journal of Cognitive Sciences & Human Development*, 2(2).
- Katz, J. (2012). *Designing information human factors and common sense in information design*. Hoboken, N.J.: John Wiley & Sons.
- Mitchell, M., & Wightman, S. (2005). *Book typography: A designer's manual*. Marlborough, Wiltshire: Libanus Press.
- Parilová, T., Mrván, F., Mizík, B., & Hladká, E. (2016). Emerging Technology Enabling Dyslexia Users To Read and Perceive Written Text Correctly. *Research in Computing Science*, 117, 173-183.
- Shneiderman, B., & Plaisant, C. (2004). *Designing the user interface: Strategies for effective human-computer interaction* (4th ed.). Boston: Pearson/Addison Wesley.
- Singhal, A. (2012). Introducing the knowledge graph: things, not strings. *Official google blog*.
- Sipe, L. R. (1998). How picture books work: A semiotically framed theory of text-picture relationships. *Children's Literature in Education*, 29(2):97–108.
- Sirkia, T. and Sorva, J. (2015). How do students use program visualizations within an interactive ebook? In *International Computing Education Research*, pages 179–188. ACM.
- Timpany, C. (2012). Designing the printed book as an interactive environment. *The International Journal of the Book*, 7(1), 11-28
- Tufte, E. R. (2006). *Beautiful evidence* (Vol. 1). Cheshire, CT: Graphics Press..
- Vanderschantz, N., Timpany, C., Whitehead, D. & Carss, W. (2010). A small-scale study into the effect that text & background colour has on processing and self-correction rates for childrens' on-screen reading. *International Journal of the Book*, 7(4), 75-88.
- Vanderschantz, N., Timpany, C., Hinze, A., & Du, Y. (2017). *Exploring Orientation Cues in both Digital and Physical Books*. In *Proc 31st British Human Computer Interaction Conference*, 10 pages, Sunderland, UK.