The Pace of Academic Prototyping: The Case of the DTOC

Dr. Stan Ruecker / Associate Professor / IIT Institute of Design
Nadine Adelaar / Researcher / University of Alberta
Dr. Susan Brown / Professor / University of Guelph
Dr. Teresa Dobson / Associate Professor / University of British Columbia
Dr. Ruth Knechtel / Researcher / University of Alberta
Andrew MacDonald / Researcher / McMaster University
Ernesto Peña / PhD student / University of British Columbia
Milena Radzikowska / Associate Professor / Mount Royal University
Geoff G. Roeder / MA Student / University of British Columbia
Dr. Stéfan Sinclair / Associate Professor / McGill University
Jennifer Windsor / MA Student / University of Alberta
and the INKE Research Group

Abstract

It is sometimes said that the academy today is undergoing rapid change, and that one of the changes we have been experiencing is in the rate of dissemination of research results. Our previous generation could comfortably spend 10 years on a project, and produce a monograph at the end, confident that they were meeting the expectations of everyone involved. Contemporary academics, on the other hand, typically publish and present on an annual cycle; research professors spending 10 years on a project without publishing would be considered to be underperforming, if not outright failing in their duties. However, in some respects the rapid cycle of dissemination masks the actual pace of academic understanding: people still need time to work, to write, to reflect. In this paper, we argue that researchers who routinely design, prototype, and test experimental systems, far from being the exception to the rule, may actually require more time to build and make their cases. As an example, we discuss the Dynamic Table of Contexts, first conceived of in 2003 and described as a concept at a DH conference in 2005. Ten years later, in 2013, it is still moving forward as a constellation of design, programming, and testing, but more importantly, as a locus of ideas about what it means to remediate a venerable print tradition in such a way as to make it acceptable to print readers while at the same time increasing its affordances.

Introduction

It is commonly recognized that academics in 2013 tend to produce more publications more quickly than was the case for academics in, for instance, 1953. Part of the reason is the current state of the academy, especially in the arts and humanities, where competition for academic jobs has resulted in earlier and more aggressive publication. It is not uncommon to see people applying for junior positions with a dozen or more papers in good journals, and in some cases a book or more. The funding agencies also have a role to play, with an emphasis on evidence of dissemination. Finally, there continue to be changes in publication models, which can provide more opportunities for a wider range of publication, although alas at this point we still have no way to recognize experimental prototypes directly through a process of peer review.

What we typically have instead is a situation where people interested in working with experimental prototypes will propose some relevant research questions, often involving user study, then write and publish about the design ideas and the user tests, with the prototypes arising as a kind of side effect of the process. There have occasionally been efforts to address prototypes more directly, but what has tended to happen is that a brief discussion of the prototype turns out to be insufficient for people attempting to assess their value as instances of new knowledge, and the length of the discussion increases over time until we are back to writing a paper for evaluation rather than having the prototype directly reviewed in some way.

We have previously suggested that the review of prototypes might leverage experience in reviewing other kinds of scholarly output (Galey and Ruecker 2010), but so far, despite a growing interest, nothing of note has emerged. All of which is not to say that experimental prototypes are not of value in the pursuit of new knowledge, and in fact, it is possible to claim that a kind of intellectual

trajectory can be recognized in looking at the changes made to prototypes over time. Our claim here is that this trajectory shows reasonably clearly that at least some academic prototypes have a relatively long lifespan as prototypes, with each successive cycle of design and development extending our understanding. As an example of the process, we will look at the conception and history of a relatively simple prototype system called Dynamic Table of Contexts (DToC).

Research Question: what are the possible new affordances of digital text?
The concept of the DToC arose as one answer to a much larger question: what are the possible new affordances of digital text? In general, these involve research areas such as text analysis and visualization. In more specific terms, we have distant reading, algorithmic criticism, text visualization, and interactive visualization. Our research projects in response to this question have included:

- the Bi Sheng electronic book
- watching the script
- the simulated environment for theatre (SET)
- No One Remembers Acronyms (NORA)
- Metadata Offer New Knowledge (MONK)
- Implementing New Knowledge Environments (INKE)
- the paper drill
- citelens
- the multitouch variorum
- the text as a string of words
- Just in Time Research (JiTR)
- repetition loops
- dialR
- the novel as slot machine
- the magic circle
- the Mandala Browser
- texttiles
- bubblelines
- conversational modeling
- the reorganizable textbook

This list is significant in two ways. First, it shows the diversity of subtopics within a single research question that can be addressed through an approach that may include the design, prototyping, and testing of concepts. Second, for anyone familiar with the list, it may bring to mind the complex interactions that occur between prototypes and projects. To take one example, the magic circle (Fig. 1) is a visualization that shows a breakdown of various contributions to some sort of whole, provided by a number of designated parts of that whole. It was originally sketched as a possible component of the MONK project, where it was going to serve as a way of showing how vocabulary use varied across different works, either by the same author or else across multiple authors. By extension, it could also be used to look at varying use of lemmatized vocabulary, stemmed vocabulary, or even parts of speech.

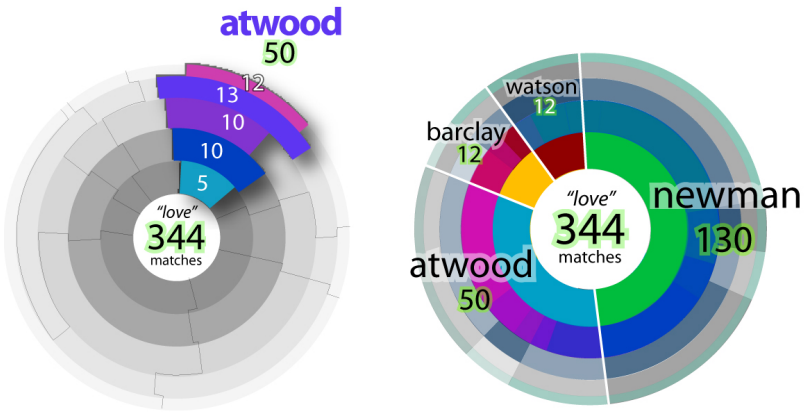

Fig. 1: The original magic circle design, showing search results for the word "love" across multiple books by multiple authors.

However, the magic circle did not find its way into the final production of MONK, and was shelved for a couple of years. It next emerged in a collaborative project on wiki authorship, where the team was looking for some means of conveying the results of an algorithm that could give people credit for their contributions to wiki pages. Somewhat at odds with the typically anonymous approach to wiki authorship, the magic circle was an innovation intended to accommodate writing environments like industry and the classroom, where anonymous contribution is deprecated in favour of knowing who has been working on what (Arazy et al. 2010).

Taking it one step further, the magic circle was included as an example of the kinds of output that could be anticipated from the interface design research team in the Implementing New Knowledge Environments (INKE) project. In this context, our intention was to apply it more broadly to questions of co-authorship, not just on wikis, but also across a number of media.

Finally, extending beyond INKE, the magic circle was adopted for production by the Canadian Writing and Research Collaboratory (CWRC), where it is intended for use in helping to manage and to provide appropriate credit for collaborative authorship across a range of projects. This one fairly small visual idea has therefore served as a component of three major projects, helping researchers to consider some of the complexities inherent in the activities they are facilitating.

A Brief Visual History of the Dynamic Table of Contexts

In the case of the DToC, we were interested in the design of the electronic book, and how the people who want to read and study books might have their experience improved over what is possible with a print book. Our contention was not that digital books would replace print books, since we had previously carried out a study that suggested that they would not, at least for dedicated readers in this lifetime (Ruecker 2002). Instead, we wanted to identify growth points, where what we would now call the skeuomorphic design, and what we then spoke of as the remediation of the print book into digital form, could be extended in ways that would be beneficial to the reader.

Our initial concept and conference paper (Ruecker 2005) asked how the print table of contents might serve as a rich-prospect browser. We identified which of the principles of rich-prospect browsing it met, and which it failed to meet, then posited a digital system that would meet all of the criteria. The purpose of meeting these criteria was that the TOC could then become a more robust tool for researchers. Our example used the dynamic insertion by the reader into the TOC of characters, dialog, and locations in *The History of Tom Jones, a Foundling*. At this stage, we were looking at some early sketches (Fig. 2) but had not begun to think about the details of programming a working prototype.

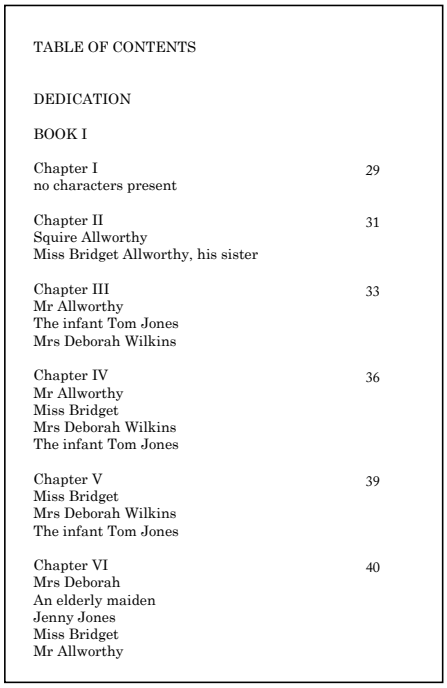

TABLE OF CONTENTS

Fig. 2: Characters in the first few chapters of *Tom Jones* • where in the book do they appear? • what are they called? • what are they doing? • how are they characterized? • what is their ontological status? (c.f. Willard McCarty on the *Metamorphosis*)

This might be an appropriate place to mention one of the virtues of the digital humanities as a field: we are at least as interested in the ideas behind the prototype as in the prototype itself. In practical terms, this means that we can discuss design concepts before we build them, allowing us to put extensive work into the design thinking rather than needing to push forward quickly into a software version. It is possible in DH to present and publish on ideas, sketches, interactive sketches, prototypes, and development systems, as well as user studies at practically every stage.

The next phase of the DToC took place in association with the Orlando Project, which is an online history of women writers in the British Isles, developed at the University of Alberta and the University of Guelph, and first published by Cambridge University Press in 2006. Our question in this case was whether an interactive TOC could be embedded somehow in the existing Orlando interface to accommodate more conventional period histories (Fig. 3).

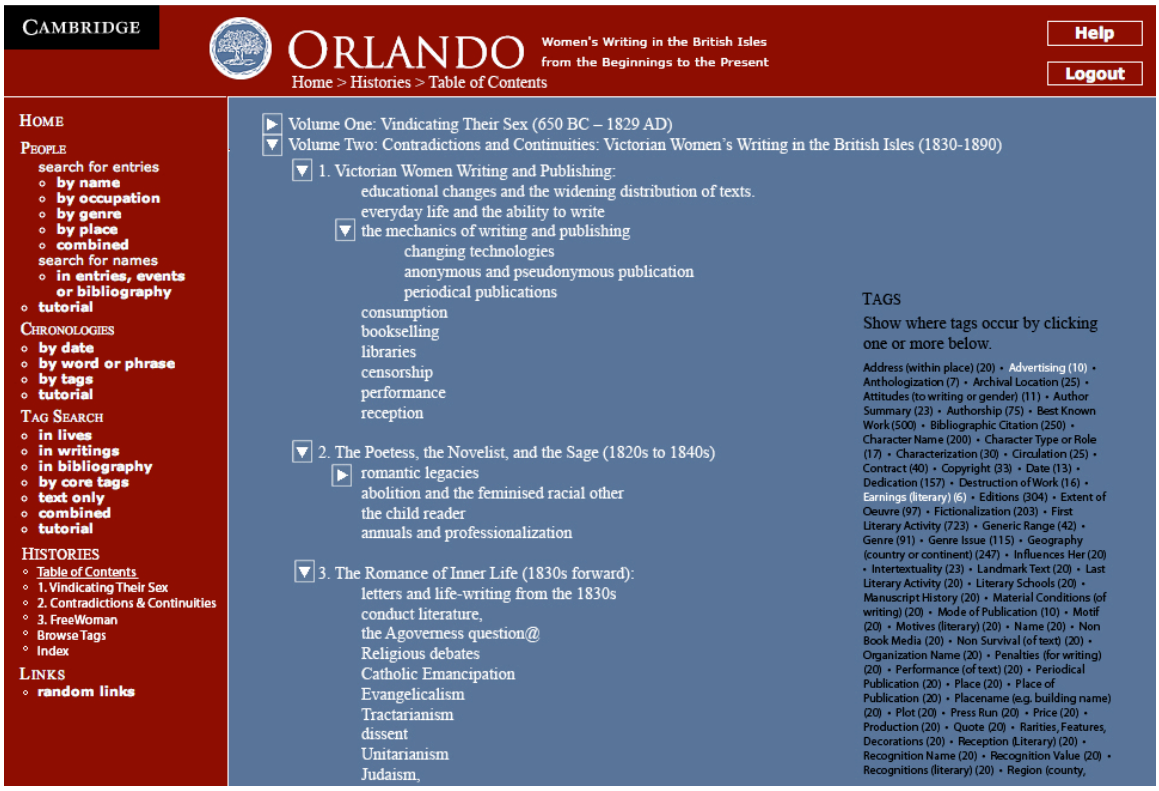

Fig 3: An interactive TOC for the Orlando Project, showing a collapsed first volume and expanded volume 2.

Since there were three proposed books, the combined tables of contents in the Orlando site would be quite complicated. We therefore added the possibility for the reader to collapse or expand selected sections of the TOC. It also became clear from the pattern of use of Orlando that it would be useful to accommodate not only XML-encoded items but also free-text searches.

An additional complication in this case was that the TOC would be situated in the current interface. It would therefore become another option in the panel on the left (shown as the first selection under "Histories". Other options related to the proposed volumes would include direct reading access to any of the three books, a display showing just their XML tags, or a conventional index. Note that in this version of the design, the tags and index were assumed to be aggregated for all three volumes, rather than broken out by volume.

Placing the history volumes as an extra item in the navigation panel is one possible strategy, based in part on the principle that readers would want to have all the navigation items available at all times. However, this approach was also potentially confusing for readers, in that the existing four options in the navigation panel (i.e. People, Chronologies, Tag Search, and Links) provided access to the biographical material in the project, rather than to the more conventional historical volumes.

At this stage, we had done some preliminary thinking about how the mechanism of adding and subtracting material to the TOC would actually look and work. The reader would be provided with a list of possible items to include, and could toggle them off and on. We recognized that we could use the XML encoding, and leverage the human-readable versions of the tags that were already in place for the Tag Search option. We tried to condense the list as much as possible by using bullets as separators, since there could conceivably be on the order of 200 tags to consider. Since the contents would already be indexed, we were also able to show a count next to each tag, indicating how many instances would be inserted when the reader toggled that tag.

These designs combined existing Orlando material with the prospective material about the proposed volumes. In practice, the tagset for the volumes would most likely differ in some respects from the tagset previously used for the biographies. However, since the sketches were based on the existing Orlando tagset, and at this point that tagset had not been publicly released, we removed it from the published version of the design (Ruecker et al. 2007).

Working with Orlando therefore provided us with a variety of ideas about ways in which a relatively simple original concept could be complicated by embedding it within a specific pre-existing interface. This led us in the next iteration to ask how the system could be designed to act in a more general way as a standalone interface for any digital text contents (Fig. 3).

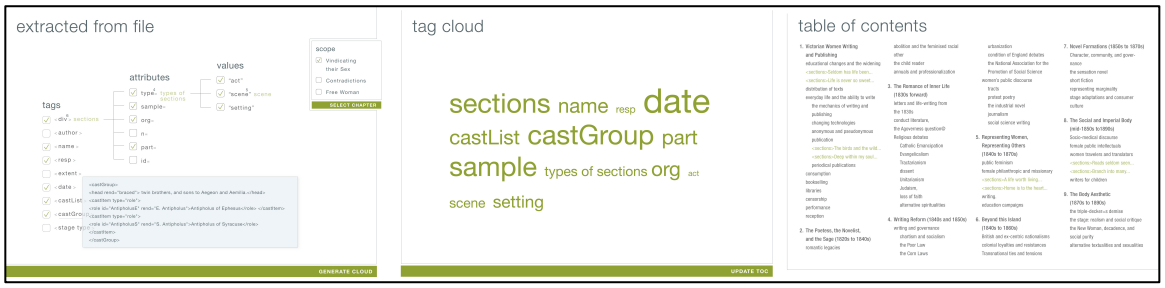

Fig 3: Design sketches by Milena Radzikowska for a generic DToC. The process goes from choices made on either of the two panels to the left and results in a display on the right, which would then be used to access the text itself.

In the generic design, the reader has the opportunity to choose from two different interfaces in order to select the items to be inserted in the TOC. The list view appears on the left; in the center is a tag cloud. Clicking on an item in either would insert the instances of that tag into the TOC on the right. Clicking on an item from the TOC would take the reader into the text.

A significant addition to the concept in this iteration is the inclusion on the left of more than just a list of tags. There is a Scope panel that allows the reader to choose which chapter to display. There is also, associated with each tag, a list of its attributes, and associated with each attribute, the values it contains. With this mechanism, the reader is provided with a finer grain of selections for which items to insert in the TOC. Since it also quite complicated, choices made in this panel can also be used to generate a much simpler tag cloud.

We subsequently had the opportunity to have the generic design built into a prototype by a group of senior undergraduate computing science students[1], who chose this project for a team programming class (Fig. 4). Their primary interest was in figuring out how to provide the necessary functionality, so attention to the design was left for another phase. This prototype was built using Google Web Toolkit (GWT).

[1] The students were Mark Bieber, Jamie Czerwinski, Xuefeng Ding, Matt Gooding, and Mike Packer. The instructor was Dr. Ken Wong.

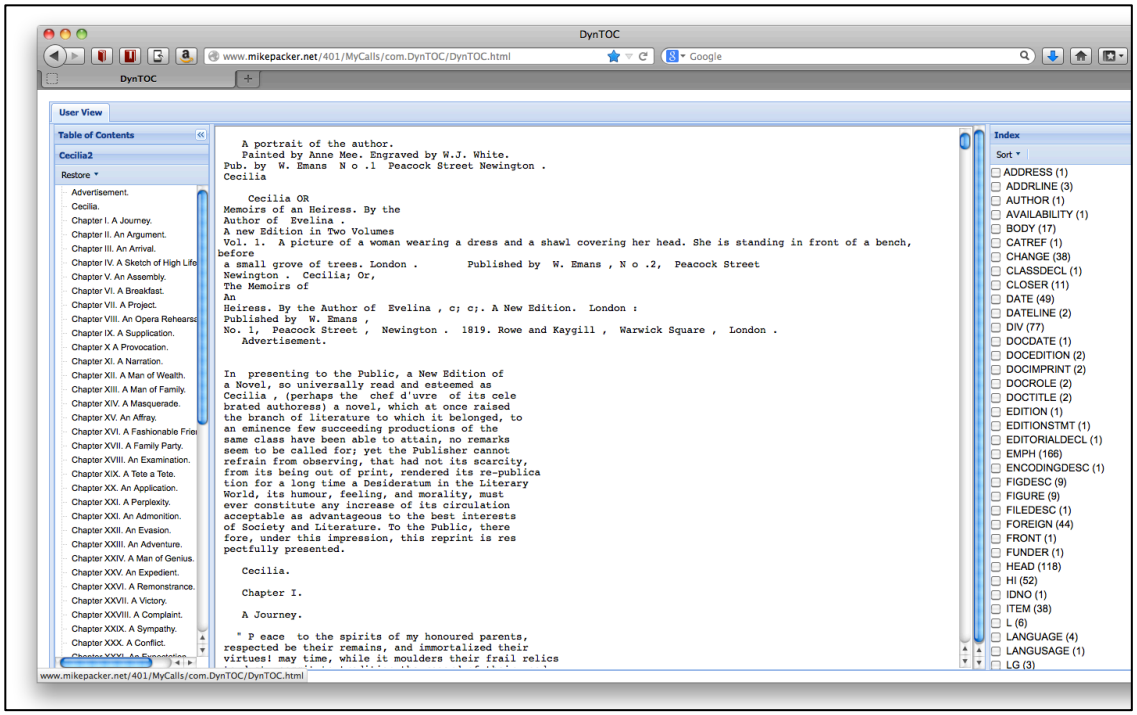

Fig 4: First prototype programmed in GWT by senior students in computing science at the University of Alberta.

The prototype was useful in examining how the TOC interacted with the XML-encoding. At this point, it became clear that we would need a curator mode, where someone would manage the collection, choose which tags should be displayed to the reader, and rename those tags to make them understandable by people who were not involved in the encoding project.

Our XML-encoded content for this prototype was Frances Burney's novel *Cecilia*, provided by the Brown Women Writer's Project. In conversations with Syd Bauman, it became clear that a useful feature to provide for the curator would be the ability to apply the same choices of tags and human-readable tagnames to any documents that shared a tagset.

Since *Cecilia* has more chapters than could appear in a list on a conventional screen at the time, inserting material of any length quickly reduced the number of chapters that were visible. We therefore inserted only a small amount of text, but added a popup that would appear on mouseover. The popup showed a reasonably large section of the text containing the tagged material.

To address the same issue of screen real estate, readers also had the option of removing chapters from the display; they could be added back with a "restore" switch at the top of the panel.

We also became aware at this stage that people may want to insert more than one tag at a time. A consequence of providing this affordance is that they would need to be able to keep track of which tag selected from the panel on the right had resulted in the contents being inserted on the left. Our solution was to preface the inserted text with the name of the tag.

Although we now had a working prototype, after the class was over we had no ready means of modifying the interface. In talking further with one of the student programmers (Mark Bieber), we also became convinced that rewriting it from scratch would provide significant advantages.

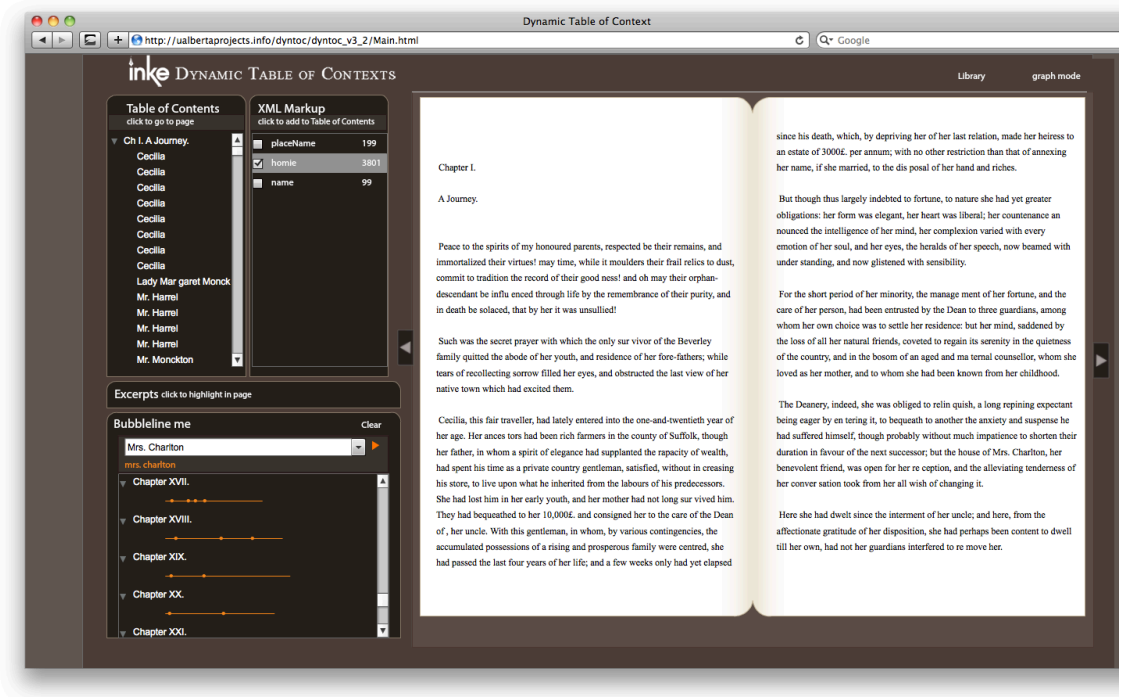

Fig. 5: Second prototype programmed in Flash by Mark Bieber; design by Milena Radzikowska.

Our opportunity to attempt this next iteration arose in connection with the INKE project, where our mandate was to look at improved interfaces for people working with electronic text. As opposed to the previous goal of creating a standalone reader with improved navigation, in this phase we were beginning to consider how the reader might interact with other features of an entire "new knowledge environment".

Working again with *Cecilia*, we produced a second prototype that incorporated facing pages in the reading panel, as well as additional text analysis and visualization tools. In particular, we added a search function, as well as a Bubblelines visualization for comparing either tagging or search results across multiple chapters. It was also possible to swap out the reading panel for a search term frequency graph.

This was the first prototype where we had the opportunity to run a user study (Dobson et al. 2012). Among other things, we learned that it would be useful to carry out user studies earlier in the process, preferably using sketches, so that we could more quickly narrow our design target before programming. For instance, our assumption that facing pages would be a skeuomorphic feature of importance to readers, the user study suggested that it was more important to accommodate people with smaller screens. Another important finding was that people were familiar with the idea that any encoding is a kind of interpretation, and they were interested in understanding more about the rationale behind the available tags they were suddenly using to navigate the text.

Simultaneously with the user study, we began to involve researchers who had been working in the history of book design, thinking about ways in which print experiments might be brought back into play within a digital environment (e.g. Nelson et al. 2011). The second prototype was also polished enough that it could be used in conversation with potential INKE project partners to suggest the usefulness and practicality of the approach.

Building on the insights from Nelson et al. and the findings of the study by Dobson et al, we therefore began a third prototype (Fig. 5), in conjunction this time with two partners: the Canadian Writing Research Collaboratory (CWRC) and the University of Alberta Press. This third prototype was also

built using the Voyant tools platform, meaning that it could potentially be robust enough to work as a development version for release to other researchers and projects. Close integration with Voyant also meant that we could conceivably find ways to leverage the many existing text analysis and visualization tools to work in conjunction with the DTOC.

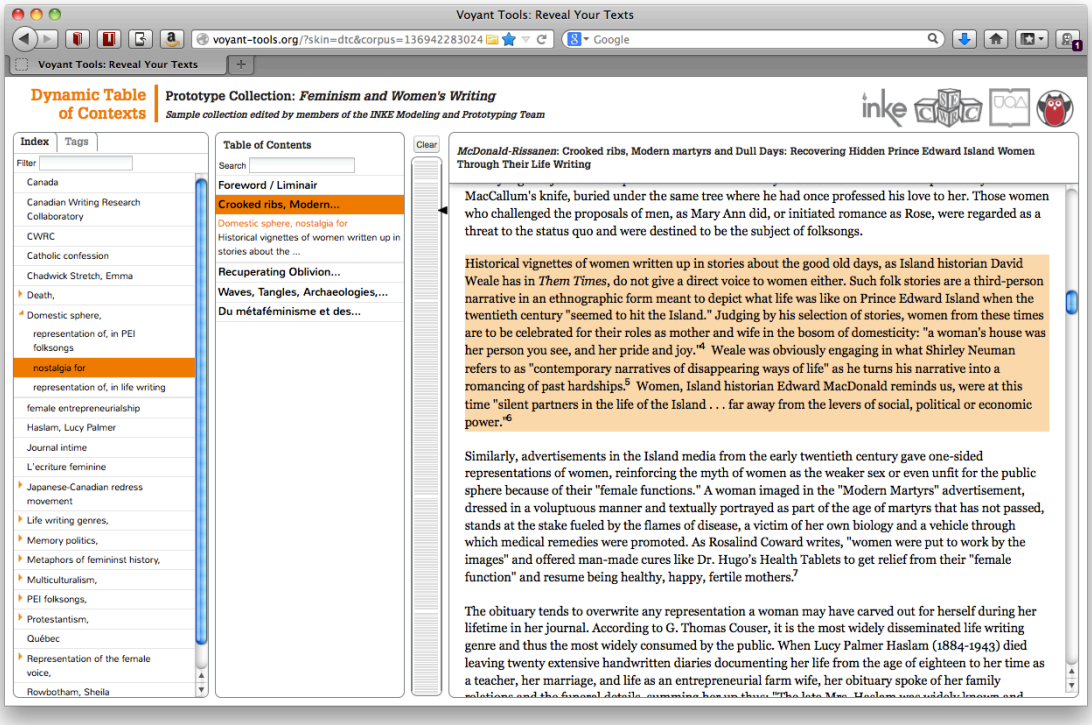

Fig 6: Third prototype, programmed by Andrew MacDonald for CWRC in Voyant using HTML5; design by Jennifer Windsor.

This version of the system adds several significant features. First is the index panel, which parallels the tag panel and serves a similar function with dissimilar contents (Brown et al. 2013). The traditional index is intended to help the reader locate specific information that in many cases will only appear once or twice within a book. Semantic encoding, on the other hand, has tended to be applied to larger concepts that occur frequently enough to make including them in the tagset seem worthwhile. It would be normal, for example, to expect to see semantic tags for material such as people's names, significant dates, and geographical locations. In an index, however, what would typically appear are the actual names, dates, and places. There is no easy mechanism for finding in an index all of the names, dates, or places at once; using an encoded file, however, will allow the reader to find all of them easily by searching for the presence of the appropriate tag.

The second addition to this prototype is the document model, which appears as a thin column of lines between the navigation panels and the reading panel. This model provides a prospect view on the entire document (Ruecker et al. 2005b), allowing the reader to see at a glance where the currently selected tags, index items, or search terms occur, since they appear as coloured lines at approximately the right position. The lines are also interactive, so they can be used to jump to the next instance in the text.

The document we were using for the third prototype is a frankentext that combines book chapters out of a recent edited collection with other book chapters that have not yet appeared in print, but have already been subjected to encoding and indexing. One advantage of this content is that it also contained footnotes, reminding us to accommodate them in the reading panel with popup balloons.

The balloons in turn required some careful thinking, since we wanted the contents to be available to cut and paste, while at the same time we wanted to avoid having them clutter up the reading panel. Our solution was to make them appear when the reader rolls the cursor over the footnote number in the text, but disappear when the cursor leaves the number—unless the user clicks on the number, at which point the footnote bubble becomes persistent until closed by clicking on a standard X in the top right corner.

Fig. 7: Voyant DTOC with panes collapsed for reading on an iPad.

Given the current prevalence of smaller reading devices such as the iPad and Kindle, the question also arose as to how difficult it might be to adjust the Voyant DTOC so that it can serve with the smaller screen real estate. As it turns out, since the panes are collapsible, it is possible to set the navigation panels to the side while reading (Fig. 7), then expand them when necessary.

Conclusions

The DTOC is one example of many of the perhaps surprisingly slow pace at which an academic prototype can proceed. It has been an intermittent focus of attention for a dozen or more researchers over a ten-year period. It is of course possible to interrupt a research trajectory involving a series of prototypes. In fact, given the exigencies of academic research funding, it is sometimes the case that a single prototype is all that is possible.

However, provided that a trajectory can be followed, it can be argued that pursuing it over an extended period, perhaps as in this case with a changing constellation of researchers, can result in an increasing degree of understanding of a variety of topics that are brought into focus by thinking through prototyping. In addition to the logistics of assembling the necessary people and resources, what the pace allows is time for reflection.

Perhaps taking a few liberties with the details of the timeline, we have essentially addressed the following sequence of questions:

2005: how does a table of contents need to be augmented to serve as a rich-prospect browser?
2006: what issues may arise in attempting to embed a table of contexts in an existing project interface?
2007: how can an electronic reading environment that includes a dynamic table of contexts be designed to act in a more general way as a standalone interface for any XML-encoded texts?
2008: what other text analysis and visualization tools might be useful to associate with the DTOC?
2009: what can we learn from a user study of a standalone DTOC prototype?
2010: what aspects of particularly interesting print TOCs and Indices can inform the DTOC?
2011: how can the DTOC be modified to accommodate scholarly editions as opposed to normal monographs or collections?
2012: using the DTOC as a case study, what can we learn by comparing semantic XML encoding with a conventional book index?
2013: what changes would need to be made to the dynamic table of contexts in order for it to potentially work on a smaller reading device such as an iPad or Kindle?

The immediate future of the DTOC will see further developments of it as a reading environment that supports other forms of visualization and text analysis. We would like it to easily access texts from a content management system. We are wondering if it would be useful for the reader to be able to dynamically reorganize the TOC panel, and whether editors would want to provide more than one form of TOC if the system were to make it possible. Then if multiple TOCs are available, should more than one be visible at a time? We have not even begun to consider what it might mean to accommodate other media, such as images or videos, but dealing with them seems like a logical next step. No doubt other possibilities will present themselves, as we ponder, build, and test our way into the future.

References

Arazy, Ofer, Eleni Stroulia, Stan Ruecker, Cristina Arias, Carlos Fiorentino, Veselin Ganev, and Terence Yau. "Recognizing Contributions in Wikis: Authorship Categories, Algorithms, and Visualizations." *Journal of the American Society for Information Science and Technology (JASIST). 61*(6), 1166-1179. 2010.

Brown, Susan, Brent Nelson, Stan Ruecker, Stéfan Sinclair, Nadine Adelaar, Ruth Knechtel, Jennifer Windsor and the INKE Research Group. "Text Encoding, the Index, and the Dynamic Table of Contexts." Paper presented at the annual Digital Humanities conference (DH2013), Lincoln, Nebraska. July 16-19, 2013.

Dobson, Teresa, Brooke Heller, Stan Ruecker, Milena Radzikowska, Mark Bieber, Susan Brown, and the INKE Research Group. "The Dynamic Table of Contexts: User Experience and Future Directions." Paper presented in the panel "Designing Interactive Reading Environments for the Online Scholarly Edition" at the DH2012 conference, Hamburg, Germany. July 16-20, 2012.

Galey, Alan, Stan Ruecker, and the INKE Research Group. "How a Prototype Argues." *Literary and Linguistic Computing*. *25*(3). 2010.

Nelson, Brent, Stan Ruecker, Milena Radzikowska, Stéfan Sinclair, Susan Brown, Mark Bieber, and the INKE Research Group. "A Short History and Demonstration of the Dynamic Table of Contexts." Paper presented at INKE: Research Foundations for Understanding Books and Reading in a Digital Age: Text and Beyond. November 18, 2011. Ritsumeikan University, Kyoto, Japan.

Ruecker, Stan. "Carrying the pleasure of books into the design of the electronic book." In Pleasure with Products: Beyond Usability. William S. Green and Patrick W. Jordan, Eds. London: Taylor and Francis, 2002. pp. 135-146.

Ruecker, Stan. "The Electronic Book Table of Contents as a Research Tool." Paper presented at the President's Panel. Congress of the Humanities and Social Sciences: Consortium for

Computers in the Humanities / Consortium pour Ordinateurs en Sciences Humaines (COCH/COSH) Annual Conference. London, Ontario. May 27- 30, 2005a.

Ruecker, Stan, Eric Homich and Stéfan Sinclair. "Multi-level Document Visualization." *Visible Language*. 39(1), 33-41. Jan/Feb 2005b.

Ruecker, Stan and the INKE Research Group. "Introducing the Dynamic Table of Contexts for Scholarly Editions." Paper presented at the Modern Language Association (MLA) Conference. Los Angeles, CA. Jan 6-8, 2011.

Ruecker, Stan, Milena Radzikowska, Susan Brown, Thomas M. Nelson, Isobel Grundy, Patricia Clements, Sharon Balasz, Jeff Antoniuk, and Stéfan Sinclair. "The Dynamic Table of Contents: Extending a Venerable List in a Digital Context." *The Potential and Limitations of a List: an International Transdisciplinary Workshop.* At the Center for Theoretical Study, the Charles University and the Philosophical Institute of the Academy of the Sciences of the Czech Republic, Prague. Nov 8-9, 2007.

Ruecker, Stan, Susan Brown, Milena Radzikowska, Stéfan Sinclair, Thomas M. Nelson, Patricia Clements, Isobel Grundy, Sharon Balasz, and Jeff Antoniuk. "The Table of Contexts: A Dynamic Browsing Tool for Digitally Encoded Texts." In *The Charm of a List: From the Sumerians to Computerised Data Processing.* Ed. Lucie Dolezalova. Cambridge: Cambridge Scholars Publishing, 2009. pp. 177-187.