

Using DocBook to Produce a Polyvalent Academic Work

Murray Maloney

Muzmo

Robert J. Glushko

UC Berkeley

R. Alexander Milowski

UC Berkeley

Abstract

Creating many different versions or customizations by configuring a collection of components is a desirable goal in many domains. A single automobile production line can support the assembly of customized variations of a car model. Software product line engineering enables the creation of many similar software systems from a shared set of software assets. In this article we discuss how a collection of content elements can create a family of related texts whose different members are generated according to configurations of variables found in the content markup. This markup is created by the author, but anyone can create a particular edition of the text by defining a configuration file at book-building time, and a reader can do this interactively at reading-time by making selections from a configuration control widget. We call this configurable collection of content elements a "polyvalent" document: "Poly" means "more than one" or "many" - "valent" means "having combining power."

There are some common challenges in all of these domains. The first is to distinguish the components that are contained in every manifestation, typically called the core, base, or platform, from those that vary, typically called the features, options, or supplements. The second is to organize the variable components to indicate the different customizations, versions, or editions that can be built by selectively combining optional components with the required ones. The third challenge is to convey to the builders, users, or others who want to use the variable components any dependencies or constraints that might exist, since not every possible combination will be feasible or sensible.

In this paper, we examine the facility with which DocBook was coerced into supporting a polyvalent text, and the challenges encountered. We observe the parallels and disjunctions among the vocabularies used in book production, the suitability of XHTML and CSS as content delivery agents, the varying

capabilities of current ePub3 readers, and the suitability of relying upon CSS and JavaScript in an ePub context.

1. Introduction

This article describes the challenges and insights that emerged in the design, development, and delivery of a book titled *The Discipline of Organizing* ([9]). TDO proposes a transdisciplinary synthesis of ideas from library and information science, computer science, informatics, cognitive science, business, and other disciplines that arrange collections of resources to enable interactions with them.¹

Organizing is a fundamental issue in many professional fields. However, these fields have only limited agreement in how they approach problems of organizing and in what they seek as their solutions. Nevertheless, despite their obvious differences, the books in libraries, the animals in zoos, weather observations in a data repository, and digital songs on a music player are all "resources" – "things that have value that can support goal-oriented activity" – that have been intentionally selected and organized. Similarly, despite their obvious differences, libraries, zoos, data repositories, and music collections can all be described as "organizing systems" – each is "an intentionally arranged collection of resources and the interactions they support."

A discipline of organizing complements the conventional disciplinary focus on specific resource and collection types (libraries organize books, museums organize art, business systems organize product and customer information) with a framework that views organizing systems as existing in a multi-dimensional design space in which different types of resources can be considered simultaneously, better exposing the relationships and contrasts among them. There are five groups of design decisions, phrased in generic language to emphasize their broad applicability: What is being organized? Why? How much? When? By what means?

A book with the ambitious goal of defining a new discipline must be broad enough to include all the disciplines that contribute to the "transdiscipline" that emerges at their intersection. It must treat each contributing discipline with enough depth so that the new concepts of the emergent discipline can be re-applied meaningfully to discipline-specific concepts and examples.

To make TDO both broad and deep without making it bloated and hard to understand required some innovations in book design and implementation. Our key idea was to tag the book's content by discipline, effectively creating a family of re-

¹MIT Press published *The Discipline of Organizing* in print and ebook formats in 2013. The published book names seventeen co-authors, led by the second author of this article, who also edited the book. The first author of this article, one of TDO's six principal co-authors, also served as the markup and production editor. The third author of this article implemented interactivity in the ePub editions. O'Reilly Media published two enhanced ebook editions in August 2014 to take greater advantage of the capabilities of the digital medium and to make it suitable as a textbook for a more diverse set of courses.

lated texts in which varying configurations or subsets of content tailor the book for different courses and perspectives. We are just beginning to explore the complex design space and tradeoffs between author, instructor, and reader contexts and the corresponding user interfaces that are best suited to any particular design.

1.1. Content Structure

Many books, especially technical and professional ones, are designed with a core body of content that is augmented by supplemental content of various types. The types of supplemental content, the structures that organize it in books, and its presentation and formatting are highly conventional (see [11], [14] for historical perspectives, [8], [20] for design guidance).

Tables, figures, illustrations, and sidebars are often supplemental content, and are usually constrained to appear as close as possible to the core text that mentions them. These types of supplemental content are usually created by the author or by people who are following the author's specifications.

Footnotes, endnotes, annotations, bibliographic citations, glossary entries, and indexes are types of supplemental content that are also closely anchored to particular parts of the core text. Footnotes and annotations are usually constrained to appear on the same page as their text anchor, but the other types of content are more typically arranged at the end of larger text units like chapters or at the end of the book. The author does not typically create some of these types of content, especially indexes.

Appendixes, commentaries, reviews, and case studies are types of supplemental content that are typically associated more coarsely with the book as a whole. People other than the book author also commonly create them.

The basic contrast between core and supplemental content is a very old one but the emergence of digital documents has enabled some new variations. Selectable links that "transport" the reader to the linked content or that "transclude" the content into the core text stream were foundational concepts of hypertext proposed by ([17]). The now familiar idea of web browser "plug-ins" or "extensions" for embedding new digital format types into documents was anticipated by ([18]) who developed "multivalent documents" and an extensible reading application in which new layers of content and their specialized interactions and behaviors could be overlaid on the "base" layer. Contemporary examples include the Hypothes.is² open annotation platform and the Lens viewer for scientific publications that allow readers to rearrange and focus on different parts of the article.

The goal of interdisciplinary comprehensiveness was undermining the coherence and comprehensibility of the TDO manuscript. At the same time in late 2011 many of the co-authors moved on to other jobs and projects, leaving a much smaller au-

² <http://hypothes.is>

thoring group led by the first two authors of this article to finish the book. This gave us an opportunity to rethink and revise the book from end to end and to attack rather than surrender to the breadth vs. depth challenges.

We decided to restructure the book to emphasize the transdisciplinary core of the new discipline of organizing while preserving the disciplinary identity of the concepts, methods, technology, and people that contributed to it. We did this by editing each chapter to more tightly focus on transdisciplinary content, extracting discipline-specific content into paragraph size chunks collected into a set of endnotes at the end of each chapter.

Some of this restructuring was straightforward because it was simply making explicit the organization of chapters and section. Many followed the "hourglass" or "inverted pyramid" organization typical of news stories and textbooks to begin with an introduction and easily understood or generic examples, followed by additional refining concepts and more specific examples. Often the last paragraph of a section contained the most discipline or industry-specific content, which we then moved into an endnote.

However, most of the restructuring required more thoughtful analysis to determine whether a paragraph should be considered core or supplemental. We considered using text processing tools to calculate term frequency statistics to locate paragraphs that contained concentrations of discipline-specific vocabulary, but these were unlikely to work well given the small size of the text units we sought to restructure. We instead used simpler heuristics that keyed on the occurrence of obscure words or proper nouns like "fonds" or "Sarbanes-Oxley" as indicators of disciplinary-specific paragraphs. Nevertheless, we discovered that we often referred to discipline-specific vocabulary when we proposed design patterns or introduced more abstract terms, so removing their first occurrences from the core text would be a mistake. We sometimes needed to replace pronouns and indirect referrals with more concrete referents to fix continuity problems caused by restructuring.

When TDO went to press in early 2013, we had factored about 24% of the chapter text into six disciplines, which within months were further refined into ten: Library and Information Science, Museums, Archives, Computing, Web, Cognitive Science, Linguistics, Philosophy, Law, and Business.

1.2. Single-source Publishing

Single-source publishing promises the ability to deliver multiple forms of a given work from a single set of source files. For our purposes, that initially meant print-ready and hypertext formats.

By late 2012, the chapters of the book, in simple Word documents, were nearing completion. The front and back matter existed only notionally in the mind of the markup editor. We wanted to publish in both print and ebook formats because we expected it would need frequent revision to stay current. We were dissatisfied with

Word as our source framework, partly because its glossary and indexing tools evaded us, and partly because of our own biases toward non-proprietary, standards-based encoding schemes. We had surveyed XML-based publishing tools and frameworks, and had narrowed our sights on either DocBook or DITA. As a matter of practicality, we chose DITA because Eliot Kimber had agreed to work with us to convert our sources and adapt the standard DITA framework to produce our required output formats. We were actively preparing to produce the work using the DITA framework in Austin, TX in early December 2012. A chance encounter at a publishing conference presented us with a different approach.

We were fortunate to become beta testers for O'Reilly's Atlas single-source publishing environment. Atlas enabled us to deliver print-ready copy, epub, and mobi versions of TDO from the same XML source files, which we marked up using the DocBook schema. DocBook contains chapters, sections, paragraphs, sidebars, lists, figures, tables, links, citations, quotes, glossary and index terms, and other elements needed for books and technical publications. Because it is straightforward to transform text with this rich markup into the other formats and assign corresponding style sheets to them, the first editions were essentially identical except for the layout and formatting flexibility, search, and hyperlinking that are intrinsic to the digital formats.

Atlas was an essential productivity boost, but we resisted the siren call of complete single-sourcing. Rather than take a break during the period that the print book was working its way through manufacturing and distribution, we decided to invest heavily in ebook-targeted content enhancements and semantic markup that were not used in the print version. Even before the print edition of TDO reached bookstores in May 2013 we were well along on an enhanced ebook design that included several dozen photos, embedded quizzes, annotation capability, and other features that took advantage of the digital reading platform.

Atlas continued to meet our requirements for producing a second printing, but DocBook, XHTML5, and ePub had all moved forward and Atlas was not keeping pace with our requirements for greater interactivity. We worked with Bob Stayton to upgrade our sources and customize the DocBook 5 framework to suit our needs. As it turns out, the people at O'Reilly Media had been thinking that DocBook was too complex for non-technical authors and were about to abandon it, so our decision turned out to be prescient in retrospect. Now that the dust is beginning to settle on their new strategy, relying upon a simpler HTML5-based source model, we find ourselves tempted again by the siren call of single-source publishing, we are tempted to consider transforming our DocBook XML to XHTML5 to rejoin the Atlas tool chain, but only if we can do it while preserving the rich semantics we have encoded in the former. (We will propose some ways we might do this in Section 3.4.)

1.3. Selective Inclusion by Discipline

Readers are familiar with the contrast between core and supplemental content and somehow decide how much of the latter to read when they encounter it in a book. However, the distinction is generally not based on disciplinary specificity, so we could not predict how it would affect the TDO reading experience.

TDO readers understood that the disciplinary labels on each endnote could help them decide whether to read it or not. We realized too late for the first print edition of TDO that we could append discipline labels to the superscripts marking the note in the core text, but we were able to do so in the ebook editions. This made it even easier for readers to be selective about supplemental content because it made it unnecessary to flip to the end of the chapter to determine the disciplinary focus of the endnote.

We informally surveyed students in courses that used TDO about how they read the book. Some students with the print edition read everything in one pass, including the endnotes by flipping back and forth to the end of the chapter. Other students using the print edition read the core text first, and then read all the notes in a second pass. Some students reported that they ignored notes that were in unfamiliar disciplines, while others said they paid more attention to disciplines that were unfamiliar.

Tracking reading behavior is technically straightforward in digital reading environments but the most common ebook readers from Apple and Amazon are proprietary and closed, and neither firm shares user experience data with authors or publishers. Software that can unzip ebooks and render them in ordinary web browsers can exploit web analytic mechanisms for tracking user events ([4]), but will students let their instructors monitor how or if they read an assigned textbook?

A more systematic survey was conducted at Berkeley in the fall semester 2014, where all students used TDO in ebook formats in a graduate course taught by the second author of this article. This survey revealed greater likelihood to read supplemental content in ebooks than with print versions, but also confirmed that both the propensity to read and disciplinary preferences for supplemental reading were highly variable. These findings gave us more motivation to find a reading-time customization approach, which we describe in Section 2.2.

1.4. Inclusion vs. Transclusion vs. Exclusion

We had initially hoped to employ transclusion as the mechanism for incorporating supplemental content; readers would be alerted to its presence with a disciplinary-specific symbol in the page margin, and selecting the symbol would seamlessly insert the content into the core text stream, perhaps subtly altering its text formatting or font to remind the reader of its supplemental role.

However, transclusion isn't supported in any existing book reader, so we were forced to rely on more traditional inclusion mechanisms of link following (and re-

turn) and pop-up notes. We preferred the latter because it better maintains the reading context, but its poor implementation forced us to rely on the former. In particular, Apple's popular iBooks reader supports pop-ups (see Figure 1), but doesn't allow link traversal from a pop-up note. Since most TDO endnotes contain citations, pop-up notes would become dead ends. We aspire to the happy compromise that presents itself in the Lucifox plug-in for Firefox (see Figure 2).

We now realize that there is a third class of design mechanisms that we need to explore. Framing our content architecture in terms of core and supplemental content assumes that readers are selectively incorporating additional content to a book. It is interesting to consider starting with the complete book and enabling readers to selectively exclude rather than include content. One possibility would be to invert the user experience we imagined for transclusion, leaving a margin symbol to indicate where the reader has chosen to exclude content. Or we might use the familiar presentation conventions of reducing the font size or graying out for content that is excluded by the filter applied by the reader.

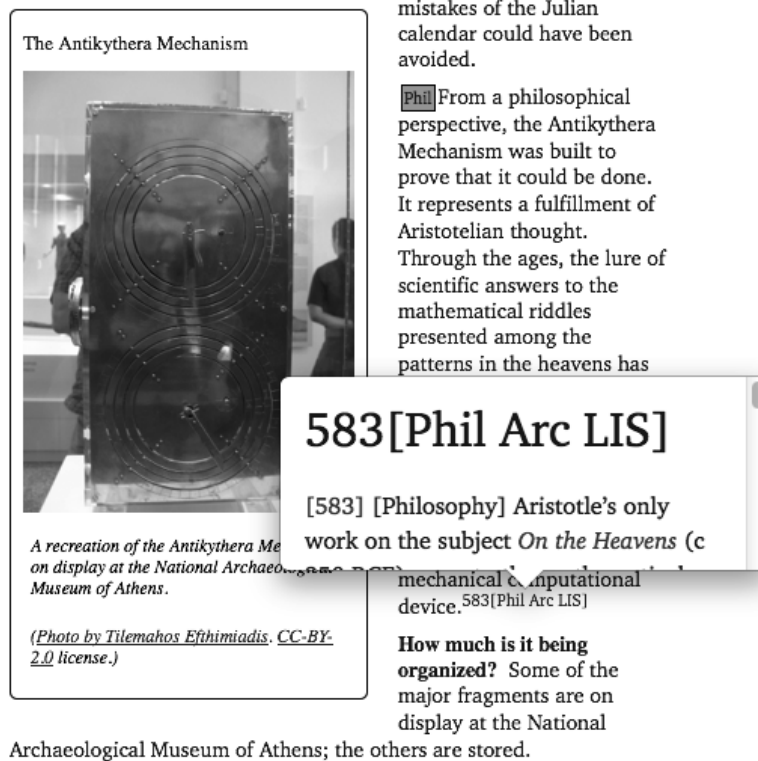
1.5. A Family of Related Books

In its first year, TDO was adopted in whole or in part by about 20 schools for a variety of courses in information organization, content management, collection development, and information architecture. However, many instructors were using only parts of TDO and asked for a simpler shorter version more suitable for undergraduate courses, which meant we needed to further refine our classification of the book's content.

We extended the idea of disciplinary labeling of content to identify an additional 15% of the chapter content as being focused on disciplinary-specific rather than transdisciplinary content. This was easier this time because it was easier to evaluate consistency and continuity when it wasn't necessary to flip back and forth between the chapter body and its endnotes to consider whether to remove a paragraph from the core.

During this time we also invested in a substantial amount of markup to enhance the amount of processable semantics represented in the source text. In addition to its structural elements, the DocBook schema contains semantic elements which we used to identify people, organizations, locations, products, applications, abbreviations, foreign phrases, and other potentially useful semantic "nuggets" that were mixed into the text. We invested in this semantic markup because we imagined being able to interconnect digital versions of TDO with other semantically described web resources by exposing it according to the conventions for "linked data" ([3], [1]) even if the current technology for ebooks was incapable of enabling it.

A more speculative form of semantic markup we explored was to label content according to its rhetorical purpose and intended audience. We classified some phrases as definitions, principles, statements, examples, parenthetical cross-refer-



The iBooks endnotes. A build-time option injects ePub3 semantics enabling an iBooks footnote interaction that presents the pop-up viewport. Sadly, it exhibits odd behavior with links. Exercising an external link presents the resource in the same cramped viewport. Exercising an internal link yields an empty viewport.

Figure 1. iBooks Popup


ences, and editorial asides; we classified some content as suitable for undergraduate, graduate, or professional audiences. We consulted some work in rhetoric, critical theory, and computational linguistics to create the former classification, but the categories and their boundaries we discovered in this diverse literature are not entirely consistent (see [2]).

We are currently exploring some different approaches for creating custom editions based on this markup. We note that there are considerable similarities between this goal of creating a "textbook family" and that of creating a "software product family" (see [12]; [13]). Furthermore, the techniques we propose have analogues with conditional compilation ([5]) and visualizations that contrast core and supplemental content ([6]; [7]).

2. A Polyvalent Academic Text

When a person or a thing is "polyvalent," it presents many different functions, forms, or facets; it is adaptable; like the "Jack of all trades". A polyvalent text is one

The Antikythera Mechanism



A recreation of the Antikythera Mechanism on display at the National Archaeological Museum of Athens.
(Photo by Tilemahos Efthimiadis. CC-BY-2.0 license.)

Why is it being organized? From a purely pragmatic perspective, the Antikythera Mechanism would have been used to accurately reckon the calendar and to navigate the globe. The Antikythera Mechanism was a relatively portable calendar, astrolabe, and global positioning, computational device, circa the 1st century BCE. With this device the mistakes of the Julian calendar could have been avoided.

Phil From a philosophical perspective, the Antikythera Mechanism was built to prove that it could be done. It represents a fulfillment of Aristotelian thought. Through the ages, the lure of scientific answers to the mathematical riddles presented among the patterns in the heavens has challenged our burgeoning intellects. The Antikythera Mechanism realized then-modern thinking on mathematics, engineering, astronomy and calendrical calculation in a portable mechanical computational device.⁵⁸³[Phil Arc LIS]

[583] [Philosophy] Aristotle's only work on the subject *On the Heavens* (c 350 BCE) avers to the mathematical symmetry and perfection in the travels of the spheres. [Archives] In 343 BCE, Aristotle was head of the Macedonian Academy, where he tutored Alexander and his future general, Soter Ptolemy. It is known that Alexander ordered his chronicler, Kallisthenes, a nephew of Aristotle, to organize the translation of all historical astronomical observations following his conquest of Babylon in 331 BCE. A transfer of scientific knowledge had begun.

[LIS] Following the deaths of Alexander and Aris[...] fabrication was not thereafter seen until the next millennium. The engineering and machining would have required trial models, accurate plans, and custom tooling. There have been various modern attempts to re-create the Antikythera Mechanism, or at least to re-create the model it seems to have manifested.⁵⁸⁴[Arc Bus Com Mus]

The Lucifox add-on for Firefox offers hypertext cues. When hovering over any link that is local to ePub3, this feature presents a tooltip with an excerpt from the target of the link and invites the reader to follow the link. This feature works with footnotes, cross-references, glossary terms, citations, and bibliography, and even for the referents of index terms

Figure 2. Hypertext cues in Lucifox

that is transformative and can adapt to the information needs of the reader. Such text often exhibits a dynamic quality where the user may express their information needs via implicit or explicit methods and the flow and contents of the text changes to represent the subset of information contained that matches their needs. Considered as a whole, the polyvalent text is a superset of all the information needs of its various consumers.

The text is written with multiple disciplines, or audiences, in mind, including footnotes and paragraphs identified as LIS, Museums, Archives, Computer Science, Linguistics, Cognitive Science, Business, Law, and so on.

Example 1. Paragraphs and Footnotes by Discipline

```
<para audience="CORE IA"><phrase role="statement">Classification makes systems more usable when it is manifested in the arrangement of resource descriptions or controls in user interface components like list boxes, tabs, buttons, function menus, and structured
```

```
lists of search results.</phrase
><footnote xml:id="endnote-394" label="394" audience="IA">
  <para audience="IA">[IA] The application of classification and
    organizing principles more generally to the design of
    user interfaces to facilitate information access, navigation,
    and use is often called <quote>Information Architecture.</quote>
    See <citation xml:id="cite_Morville2006-7.1"
      linkend="Morville2006">(Morville and Rosenfeld 2006)</citation>.
  </para></footnote></para>
```

Example 2. Bibliography by Discipline

When a bibliography entry is cited only in one disciplinary endnote, and is not mentioned in the surface text, we may choose whether to assign that entry only to the discipline specified by its parent endnote, or rather to also mention other disciplines so that the entry can appear in other editions, even though the corresponding citation and endnotes may not also appear. Here, we have an entry that will only appear in editions in which information architecture, computing, or the web are selected:

```
<biblioentry xml:id="Morville2006" audience="IA Computing Web">
  <authorgroup>
    <author><personname><firstname>Peter</firstname>
      <surname>Morville</surname></personname></author>
    <author><personname><firstname>Louis</firstname>
      <surname>Rosenfeld</surname></personname></author>
  </authorgroup>
  <pubdate>2006</pubdate>
  <title>Information Architecture for the World Wide Web</citetitle>
  <address>Sebastopol, CA</address>
  <publisher><publishername>O'Reilly</publishername></publisher></biblioentry>
```

Some structural elements of the content are identified by user levels, including Professional, Instructor, Graduate, and so on. These elements are used for the edition-specific content, such as the Stop and Think exercises in the Core Concepts Edition.

Example 3. Userlevel selection

```
<sidebar userlevel="Professor Instructor Graduate Undergraduate"
  xml:id="StopAndThink-2.4.2.1-search-engines">
  <title>Stop and Think: Browsing for Books</title>
  <para>How does the experience of browsing for books in a library
    or bookstore compare with browsing using a search engine?
    What aspects are the same or analogous in all these contexts?
    What aspects are different?</para></sidebar>
```

Some structural elements are identified by vocation, including Archivist, Curator, Linguist, and so on. These are used for edition-specific content, such as the cover, subtitle, ISBN, and so on.

Example 4. Edition-specific

```
<subtitle userlevel="Professional"
  audience="LIS Archives Museums
    Computing IA Web Law Business
    CogSci Linguistics Philosophy"
>Professional Edition</subtitle>

<subtitle userlevel="Undergraduate"
  audience="CORE"
>Core Concepts Edition</subtitle>

<subtitle userlevel="Librarian Archivist Curator"
  audience="LIS Archives Museums"
>Academic Edition: Memory Institutions</subtitle>

<subtitle userlevel="Philosopher Linguist"
  audience="CogSci Linguistics Philosophy"
>Academic Edition: Sensemaking</subtitle>
```

2.1. Static Multivalent Editions

We modified the build process we used to produce the initial ebooks to produce an ebook that contains whatever set of disciplines is specified in a configuration file. An XSLT script filters content based on the values assigned to AUDIENCE and USERLEVEL attributes when we assemble books from our DocBook XML source files.

This approach turns the source files for a book into a family of books with a common core extended with discipline-specific content. With many disciplines the combinatorial possibilities make this an extremely large family (with eleven discipline types, there are 2048 distinct combinations of zero to eleven), and even if we apply strong reasonableness or familiarity constraints it is still easy to imagine many subset configurations of disciplines that would generate appealing custom textbooks:

- Academic Edition: Memory Institutions (LIS, Museums, Archives)
- Academic Edition: Sensemaking (Cognitive Science, Linguistics, Philosophy)
- Academic Edition: Informatics (Computing, Information Architecture, Web, Business, Law)
- Academic Edition: Information Architecture (Information Architecture, Linguistics, Web)

However, for business and marketing reasons we decided, with our publisher, to go to market with just two combinations that define the endpoints of possible disciplinary customization:

- Professional Edition (all disciplinary endnotes)
- Core Concepts Edition (no disciplinary endnotes, added Stop and Think exercises)

We expect that as more schools adopt the book for a wider range of courses and student populations we will gain experience that might cause us to market other combinations.

Static editions are customized for a particular set of audience and user-level parameters by filtering content, based on profiles of DocBook 5 attribute values. Any given content element may be a member of one or more user levels and audiences. Membership in the set is primarily dependent on any one of the @audience values being selected.

For example, two paragraphs might be tagged as @audience="Computing IA Linguistics" and @audience="Computing Business Law". Those paragraphs will be members of the set if any of those audience values are selected in the build configuration. So long as "Computing" is selected, both paragraphs will be included.

Similarly, selection of one or more user levels will restrict the set accordingly.

Example 5. Configuration Excerpt

```
<group xml:id="Professional"
  name="Professional Edition"
  ref=" professional
      core lis archives museums
      computing ia web law bus
      cogsci ling phil"/>

<group xml:id="Undergraduate"
  name="Core Concepts Edition"
  ref=" undergrad core"/>
```

Thus, we were able to produce a static "Professional Edition" alongside a static "Core Concepts Edition" for undergraduates. These editions represent the two ends of the audience spectrum: all footnotes and no footnotes. We are currently testing the production of more nuanced editions for "Memory Institutions," "Informatics," and "Sensemaking"

Example 6. Academic Editions

```
<group xml:id="Memory"
  name="Academic Edition: Memory Institutions"
  ref=" graduate librarian archivist curator
      core lis museums archives"/>
```

```
<group xml:id="Informatics"
  name="Academic Edition: Informatics"
  ref=" graduate programmer
      core computing ia web bus law "/>

<group xml:id="Sensemaking"
  name="Academic Edition: Sensemaking"
  ref=" graduate philosopher linguist psychologist
      core cogsci ling phil"/>
```

The “Instructor Edition” will include supplemental material, such as in-class exercises, essay assignments, and typical examination questions. Because userlevel and discipline are orthogonal axes of selection, we will be able to create static editions that correspond to the Academic Editions, as well as dynamic polyvalent editions.

Example 7. Instructor Edition

```
<group xml:id="Instructor"
  name="Instructor Edition"
  ref=" instructor
      core lis archives museums
      computing ia web law bus
      cogsci ling phil"/>
```

As a tool for development, two special editions were designed for editorial and production purposes. The *Editor's Edition* serves as a staging ground for new material that is not quite ready for inclusion in the main corpus, and includes notes left by the editors to remind each other about needed re-writes, citations, glossary entries, and so on. The *Markup Edition* of the work includes an extra chapter, *Production Notes*, written during development and used to test DocBook markup, transformed into XHTML, and manifested in print and ePub results, using CSS to control the presentational and behavioral characteristics.

Example 8. Editor and Markup Editions

```
<group xml:id="Editor"
  name="Editor's Edition"
  ref=" editor
      core lis archives museums
      markup production publishing
      computing ia web law bus
      cogsci ling phil"/>

<group xml:id="Markup"
  name="Markup Edition"
```

```
ref=" marker producer publisher
      core lis archives museums
      markup production publishing
      computing ia web law bus
      cogsci ling phil"/>
```

2.2. Dynamic Polyvalent Editions

A second approach to customization is motivated by the reality that not all students in a particular course have the same disciplinary backgrounds and interests, and not all parts of a book require or permit the same pattern of disciplinary coverage. In TDO, for example, Chapter 3 on "Resources" discusses philosophical topics about ontology and identity like "carving nature at its joints" and "the ship of Theseus", in contrast with Chapter 8 on "The Forms of Resource Description" that has much more need to focus on technical architecture and implementation concerns. So it would be desirable for disciplinary customization to be determined by the reader in response to his preferences given a particular type of content.

To enable "reading time" or dynamic customization, we modified the build process to convert the discipline AUDIENCE attribute into a CLASS attribute in the generated XHTML. The third author of this paper implemented a prototype interface (see Figure 3 and Figure 4) that uses CSS and JavaScript to insert a list of check boxes before each section of the book; the reader can select which disciplines to include and exclude; the ebook dynamically re-formats accordingly by modifying the CSS display property of the affected elements. For convenience, groupings of discipline types that correspond to a named edition are readily chosen from a pull-down menu; this approach allows a single product to dynamically morph itself into any of the family members. However, its reliance on JavaScript limits its deployment to a select few ebook readers.

Although this prototype script is sufficient for trivial demonstration, we have had to limit its capability and utility for lack of local storage. Ideally, we would offer a facility to set and store user preferences to adjust the presentation of semantically-significant elements, under different reading scenarios. A student reading a book for the first time may want it all, while a student preparing for an examination may prefer to see only content that relates to their field.

Example 9. Excerpt from `tdo-toggle.js`

```
TDOToggle.prototype.update = function() {
  var css = "";
  for (var key in this.checkboxes) {
    if (!this.checkboxes[key].checked) {
      css += "."+this.checkboxes[key].value+" { display: none; }\n";
    }
  }
}
```

```
this.styleElement.innerHTML = css;
}
```

The screenshot shows a web interface for editing a bibliography. At the top left, there is a 'Hide Customization' button. Below it is a 'Disciplines' section with a 'Sensemaking' dropdown menu. Underneath, there are two columns of checkboxes for selecting disciplines. The first column includes Core, LIS, Museums, Archives, Computing, Information Architecture, Web, and Cognitive Science. The second column includes Linguistics, Philosophy, Law, Business, Markup, Production, and Publishing. A 'Bibliography' section is visible below the disciplines, containing a 'NOTE' box with instructions on how entries are formatted and linked. To the right of the disciplines menu, a list of bibliographic entries is displayed, each with a citation key, author, year, title, journal, volume, issue, pages, and a link to the full text. The entries include [Atran1987], [Barsalou1983], [Battistella1996], [Bentivogli2000], [Bolshakov2004], [Boroditsky2003], [Boroditsky2010], and [Boroditsky2011].

Screenshot of Bibliography with Customize script activated and Sensemaking Edition selected.

Figure 3. Sensemaking Edition

In the second of the Bibliography screenshots, we can see that a different set of disciplines are selected and the bibliography entries correspond to the new selection. The only entries in common are [Abel2014] and [Aristotle350BC]; the first being a CORE entry and the second being unattributed to a specific discipline. The decision to not attribute [Aristotle350BC] to a specific discipline was an editorial one, purposely including the entry in every bibliography, in spite of the fact that the citation is contained within a discipline-specific endnote paragraph.

3. Reflections and Directions

Wherein we reflect upon our decisions and explore our future directions.

3.1. Core vs. Supplemental

We restructured TDO into core and supplemental content relatively late in the authoring process as a response to multidisciplinary bloat, and it was often necessary to rearrange and revise paragraphs to preserve syntactic and conceptual continuity.

Hide Customization

Disciplines

Memory Institutions

Core
 LIS
 Museums
 Archives
 Computing
 Information Architecture
 Web
 Cognitive Science

Linguistics
 Philosophy
 Law
 Business
 Markup
 Production
 Publishing

Bibliography

NOTE

In this ebook, the bibliography presents an alphabetical listing of entries which detail the authors, publication dates, titles, publishers, issues, volumes, page numbers, and digital object identifiers associated with the works cited throughout this book.

Each entry begins with an identifier based upon principal author and date of publication; in cases where a principal author is associated with multiple works in a given year, a suffix is appended to differentiate. These identifiers serve as hypertext links to guide the gentle reader back to a corresponding primary citation within the chapters of the book. -- MM

[Abel2014] Scott Abel and Rahel Anne Baillie. 2014. *The Language of Content Strategy*. XML Press. [CORE](#)

[Anderson2001a] James D. Anderson and José Pérez-Carballo. 2001. "The nature of indexing: how humans and machines analyze messages and texts for retrieval. Part I: Research, and the nature of human indexing". *Information Processing and Management*. Vol. 37. pp. 231-254. <http://www.sciencedirect.com/science/article/pii/S0306457300000261>. doi: 10.1016/S0306-4573(00)00026-1. [LIS](#)

[Anderson2001b] James D. Anderson and Jose Perez-Carballo. 2001. "The nature of indexing: how humans and machines analyze messages and texts for retrieval. Part II: Machine indexing, and the allocation of human versus machine effort". *Information Processing and Management*. Vol. 37. pp. 255-277. <http://www.sciencedirect.com/science/article/pii/S0306457300000467>. doi: 10.1016/S0306-4573(00)00046-7. [LIS](#)

[Anderson2001c] Stephen R. Anderson. 2001. "Morphology". In *The MIT Encyclopedia of the Cognitive Sciences*. Robert Wilson and Frank Keil, (Eds.). 562-563. Cambridge, MA. A Bradford Book. [LIS](#)

[Anderson2003] John Anderson. 2003. *Art Held Hostage: The Battle over the Barnes Collection*. New York. W. W. Norton & Company. [Mus](#)

[Anderson2006] James Anderson and Melissa Hoffman. 2006. "A Fully Faceted Syntax for Library of Congress Subject Headings". *Cataloging & Classification Quarterly*. Vol. 43. no. 1. pp. 7-38. [LIS](#)

[Argott2009] *The Art of the Steal*. 2009. Don Argott. New York: IFC Films. [Mus](#)

[Aristotle350BC] Aristotle. 350 BC. *On the heavens. (De Caelo et Mundo)*. Ancient Greece. <https://archive.org/details/decaeleofric00aristoft>.

[BarIlan2008] Judit Bar-Ilan. 2008. "Informetrics at the beginning of the 21st century—A review". *Journal of Informetrics*. Vol. 2. no. 1. pp. 1-52. <http://www.sciencedirect.com/science/article/pii/S1751157707000740>. doi: 10.1016/j.joi.2007.11.001. [LIS](#)

Screenshot of Bibliography with Customize script activated and Memory Institutions Edition selected.

Figure 4. Memory Institutions Edition

For example, endnotes must be anchored at the end of sentences, most often at the end of paragraphs, and generally at the end of a section.

We concluded that it would have been much easier to write a book with this core + supplement architecture if we had started with this architecture in mind. This of course is conventional wisdom in software engineering; re-factoring is harder than building in modularity from the outset on a more generic platform designed to be extended with plug-in components. Nevertheless, several TDO co-authors and instructors have proposed to add additional categories of discipline-specific content to make the book a better fit to their courses and expertise, and we were able to add the eleventh disciplinary category for "Information Architecture" relatively late in the process of creating the 2014 editions. Ideally, we can identify "discipline editors" who are responsible for their evolution, adding new endnotes, sidebars, or other supplemental content as required.

Example 10. Supplemental Content

```
<sidebar audience="IA">
  <title>Information Architecture</title>
  <para>...<footnote audience="IA">
    <para>...</para> </footnote> </para>
</sidebar>
```


3.2. User Interaction

The ePub3 platform offers many features to enable the reader to interact with an ebook and often offers the reader mechanisms by which the interface can be modified to suit the individual reader's preferences, although much depends upon the platform through which the reader interacts with the ebook. The reader typically has access to a table of contents and a search interface, and often has control over some facets of page layout. Many ePub platforms also offer helpful features that may fail to account for some sophisticated document representations.

Search facilities within appropriately encoded ePub context can and should be more context sensitive. For starters, search results that distinguish among first use, definitions, mentions in keyword metadata, mentions in text and titles, mentions in a bibliography or glossary entries, and mentions in an index.

A table of contents is typically represented as a series of nesting lists with titles and references. In print, such a list might include the titles of chapters, sections, and even sub-sections. In electronic media there is no reason to limit the level of detail that we provide in a table of contents, but neither is it necessary to always present every level of the hierarchy. The ability to control presentation of the table of contents would be especially welcome.

DocBook and ePub3 each include the concept of a bibliography, glossary, and index; XHTML and CSS do not. This disjunction between the structural and elemental components of the logical work and those of the delivered product impedes the efforts of the content creator to reliably deliver utility to the consumer. We hope to discover an appropriate set of publishing semantics to allow us to encode the structure of academic works to enable discovery and interaction with the book, as a book, and imbue its logical sub-components with familiar affordances.

3.3. XHTML5, CSS and JavaScript

We observe that there exist both parallels and disjunctions among the various book production models and vocabularies employed by XHTML5, CSS, and JavaScript in an ePub context. Mismatches among these have thwarted our attempts to produce dictionary-style headings for the bibliography, glossary, and index.

Popular ePub3 reader platforms offer nominal author- or user-control over presentation of paged content or text styling. We observe inconsistent application of XHTML and CSS features across popular reader platforms. For example, building collapsible list structures for tables of contents, bibliographies, glossary, and index could be achieved with universal support for the `xhtml5:detail` element; better yet would be providing those semantics through CSS.

Example 11. CSS Precedence

Polyvalent style sheets are a challenge because of the CSS precedence rules. Given that any element may be attributed to one or more disciplines, the order of stylesheet rules affects which actions are fired. A paragraph that is attributed to IA and Web may be labeled as one or the other, depending on the order in which they appear in the stylesheet. In this example, Web wins because it is the latest, or most recent declaration.

```
p.IA::before { content: "IA"; ... }  
  
p.Web::before { content: "Web"; ... }
```

We have observed that we cannot rely on the availability of JavaScript on all reader platforms. We expected that portable reader devices might not have JavaScript capability; we were disappointed that the Lucifox add-on for Firefox disables embedded JavaScript.

3.4. Semantic Enablement

Semantic enablement is a process that relies upon the content creator to annotate their material with actionable semantic labels, and upon a downstream processor to leverage those semantics for presentation to and interaction with people. The ePub3 platform, for example, mandates that all book archives include a table of contents that is subsequently used by the reader platform to provide a hierarchical list of hypertext links.

We have had to limit our exploitation of semantic enablement while waiting upon new versions of ePub3 and DocBook5. In the interim we have relied upon DocBook's @role and XHTML's @class and @rel to signal semantics downstream. We are now considering the application of RDFa to DocBook 5.1 sources to produce rich XHTML documents containing hybrid meta-data stores leveraging native XHTML5 semantics, supplemented with a surface layer of RDFa, and provisioned with compact JSON triple stores to support JavaScript interactions.

3.5. Distributed Authoring and Publishing

Our ultimate goal is to implement a distributed authoring and publishing system in which new content can be dynamically discovered and logically included in the family of related books. We are exploring the application of DocBook5 annotation and advanced hypertext features to enable two-way linking between TDO and content created by and for the schools in which the book is being used. We are also re-considering Atlas as single-source repository for delivery to multiple media be-

cause of the obvious benefits of relying on a generic industrial-strength platform rather than maintaining our own customized one.

4. Acknowledgments

We are grateful for technical assistance provided by Eliot Kimber, Bob Stayton, Liam Quin, George Kerscher, Håkon Lie, Adam Witwer, Nellie McKesson, Fred Chasen, and the oXygen technical support team..

Bibliography

- [1] Allemang, D., & Hendler, J. (2011). Semantic web for the working ontologist: effective modeling in RDFS and OWL. Elsevier.
- [2] Bitzer, L. F. (1992). The rhetorical situation. *Philosophy & rhetoric*, 1-14.
- [3] Bizer, C., Heath, T., & Berners-Lee, T. (2009). Linked data-the story so far. *International journal on semantic web and information systems*, 5(3), 1-22.
- [4] Chasen, F., Harnell, J., and Renold, A.J. 2014. Future Press. Unpublished Master's Project, UC Berkeley School of Information.
- [5] Couto, M. V., Valente, M. T., & Figueiredo, E. (2011, March). Extracting software product lines: A case study using conditional compilation. In *Software Maintenance and Reengineering (CSMR), 2011 15th European Conference on* (pp. 191-200). IEEE.
- [6] Feigenspan, J., Papendieck, P., Kästner, C., Frisch, M., & Dachsel, R. "FeatureCommander: colorful# ifdef world." In *Proceedings of the 15th International Software Product Line Conference, Volume 2*, p. 48. ACM, 2011.
- [7] Feigenspan, J., Kästner, C., Apel, S., Liebig, J., Schulze, M., Dachsel, R., Papendieck, M., Leich, T., & Saake, G. "Do background colors improve program comprehension in the# ifdef hell?." *Empirical Software Engineering* 18, no. 4 (2013): 699-745.
- [8] Glushko, R. J., & McGrath, T. (2005). *Document engineering*. MIT Press.
- [9] Glushko, R. J. (Editor) (2013). *The Discipline of Organizing*. MIT Press.
- [10] Glushko, R. J. (Editor) (2014). *The Discipline of Organizing*. O'Reilly Media.
- [11] Grafton, A. (1999). *The footnote: A curious history*. Harvard University Press.
- [12] Kästner, C., Apel, S., & Kuhlemann, M. (2008, May). Granularity in software product lines. In *Proceedings of the 30th international conference on Software engineering* (pp. 311-320). ACM.

- [13] Krueger, C. (2002). Easing the transition to software mass customization. In *Software Product-Family Engineering* (pp. 282-293). Springer Berlin Heidelberg.
- [14] Kilgour, F. G. (1998). *The evolution of the book*. New York and Oxford: Oxford University Press.
- [15] Klein, J. T. (2010). A taxonomy of interdisciplinarity. In R. Frodeman, J. T. Klein & C. Mitcham (Eds.), *The Oxford Handbook of Interdisciplinarity* (pp. 15-30). Oxford: Oxford University Press.
- [16] Melville, Herman. (1851). *Moby Dick, or The Whale*. London. Richard Bentley.
- [17] Nelson, Theodor Holm. "A File Structure for the Complex, the Changing and the Indeterminate." *Proceedings of the ACM 20th National Conference* (1965), pp. 84-100
- [18] Phelps, T. A., and Wilensky, R. 2000. Multivalent documents. *Communications of the ACM*, 43(6), 82-90
- [19] Teufel, S., & Moens, M. (2002). Summarizing scientific articles: experiments with relevance and rhetorical status. *Computational linguistics*, 28(4), 409-445.
- [20] Williams, R. (2005). *The Non-designers Design Book: Design and Typographic Principles for the Visual Novice*. *Non Designer's Design Book*.