DocMatrix: Self-Teaching from Multiple Sources

Elena L. Glassman
MIT CSAIL, Google, Cambridge, MA, USA
glassman@alum.mit.edu

Daniel M. Russell
Google, Mountain View, CA, USA
drussell@google.com

ABSTRACT
Learners have a wealth of online resources to help them teach themselves new knowledge. However, not all resources are of equal quality or appropriateness for a learner, given the particular set of prior knowledge they bring to the learning task. Without a teacher, finding appropriate sources that shed light on a topic, collectively or individually, is critical. We present DocMatrix, an interface for viewing multiple documents in parallel, with three key features: a grid of document viewers, a common term sidebar, and enhanced tables of contents. It is designed to let an interested learner view, filter, highlight, and search many documents on a topic simultaneously. We implemented the DocMatrix interface for Google Books, and ran a user study of the prototype. The results of this study indicate DocMatrix allowed users to find, read, and synthesize more information than a traditional single-book interface.

ACM Classification Keywords
H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION
The internet has become a vast platform for just-in-time learning. Learners can consume information from a variety of online sources in order to synthesize new knowledge about possibly complex domains, e.g., when they are considering politician's platforms during elections, when they are weighing different medical treatments, or studying a new technical domain in anticipation of a career move. Tutoring is Bloom's [2] gold standard learning environment, where the teacher customizes their presentation of information to fit the learner's prior knowledge. But in many cases, learners have no tutor. What can we build to help learners teach themselves from a diversity of online sources of varying quality and appropriateness for them?

Some people teach themselves by collecting a set of books on a topic, spreading them out on a physical desktop, and reading multiple texts in parallel. These learners may be trying to simultaneously read broadly and in-depth on a topic, even though these are contradictory goals; using text from multiple sources to illuminate each other; or looking for explanations at the right level of granularity or appropriateness for them, given their own prior knowledge and skills [24, 6]. When an explanation is confusing in one text, the reader can quickly switch to another presentation of the same topic, possibly with a clearer presentation. With multiple texts instantly at hand, a sophisticated reader can rapidly build up a deep, rich understanding of a topic area.

We present DocMatrix, an interface design for supporting self-teaching from multiple structured documents. When the user enters a search query, DocMatrix returns a collection of documents, i.e., books in Google Books, embedded into a webpage in a grid so they can be viewed *simultaneously in parallel*. A sidebar displays common terms, ordered from most to least common, that occur across all the documents' section titles in the collection. When the reader clicks a term in this sidebar, DocMatrix filters and exposes the sections within each document on that topic. We aim to make it easy for readers to consume multiple explanations and viewpoints on a common topic or idea. The name DocMatrix originally just alluded to the rectangular array of embedded documents. However, we believe now that it is an even more apt name, because, like matrix decomposition, DocMatrix allows users to break a large document corpus down into its component parts that, in varying proportions, make up each document in a collection.

We implemented a version of this design for books using the Google Books API. In this instantiation, the books' tables of contents are the section titles. We ran a user study on nine subjects and evaluated, quantitatively and qualitatively, how the interface affected users' reading behavior, compared to a standard alternative. Subjects believed that with DocMatrix, they could more easily assess the usefulness of a book and synthesize information across multiple sources. When using DocMatrix, they interacted with more books and labeled more paragraphs, definitions, examples, and charts as helpful.

The contributions of this paper are:

- A new user interface for reading structured documents in parallel, with three key features: a grid of document viewers, a common term sidebar, and enhanced tables of contents;
- A user study showing that DocMatrix allowed users to find, read, and synthesize more information than a traditional single-book interface.

ASIST 2016, October 14-18, 2016, Copenhagen, Denmark.

Figure 1. The DocMatrix interface for Google Books displays a grid of embedded books returned by a query. Here, the first six of forty books in the collection, retreived with the query *algebra*, are within the reader's view.

RELATED WORK

Synthesizing Knowledge Across Sources and Modalities

The synthesis of understanding derived from multiple sources is critical to journalism and humanities scholarship and in technical fields, like mathematics.

Humanities Scholarship and Journalistic Analysis Wineburg [24] shows how students of history come to their understanding of complex events. One important behavior is the students' use of multiple simultaneous documents to understand context. Wineburg finds that "...context is everything ...who wrote something; what their political view is; what the situation in the world is at that moment ...you need to see the situation from many points-of-view ..."

Software has recently been built to help scholars and journalists analyze and synthesize knowledge across sources. For example, the AP's Overview Project is an example of software designed to help journalists analyze thousands of documents. Similarly, WordSeer [12] allows scholars in the humanities to make sense of a corpus of relevant texts by providing the ability to look at multiple sources and do textual analysis of the content. Crowdlines [10] employed crowd-sourcing to help people learn and synthesize information from diverse online sources. Since humans are skilled at evaluating high-level structure and making connections between sources, crowdworkers created outlines for important topics that included diverse perspectives from multiple document sources.

Shahaf et al. [16] created algorithms for "information cartography," algorithmically sub-sampling large collections of documents on a common topic and laying them out as a 2-D map of interrelated documents for users to explore and read. This technique has been applied to scholarly publications and news articles on complex current events, such as the European debt crisis or the Israeli-Palestinian conflict. These domains contain interrelated parallel story-lines that evolve and intersect over time, and are represented as such in the resulting Metro Maps of information.

Science, Technology, Engineering and Math (STEM) The simple question, "What machine learning books are accessible and appropriate for my high school-aged daughter?" recently kicked off a very lively discussion on a corporate engineering mailing list. It is a question that humans with a model of the learner, e.g., high school student, and the subject matter, e.g., machine learning, can answer well. Jardine [7] generated reading lists specifically for novices hoping to become experts in a particular area, using a personalized pagerank function and Latent Topic Models.

Educational psychologists have found that multiple explanations within and across multiple modalities can help students learn mathematical problem solving. Both Tabachneck et al. [19] and Cox and Brna [4] found that students fared better at problem solving when using multiple strategies and/or representations, such as diagrams, written algebra, tables, and natural language. Ainsworth points out that giving students an opportunity to consider different representations may help them overcome the weaknesses of any particular representation.

This is also supported by authors of Metacademy.com, a popular online resource for teaching yourself machine learning: "A good general piece of advice is to consult multiple resources. Different textbooks or courses will explain something from a different perspective ...[O]ften when reading one, you get an aha! moment for something which didn't make sense in the other. Unfortunately, this option might not

be practical unless you have access to a university library." [11]

Library-Inspired Software Software has also been developed to support book users more generally. The display of physical books on a library shelf according to the Dewey Decimal system can serve, at the book shelf level, as a mechanism for finding a set of topically related books to dig into further. Virtual versions of this interaction, like Bohemian Bookshelf [20] and the Harvard Library Innovation Lab's Stack View [8], use modern information visualization techniques and metadata, such as patron usage, to help readers serendipitously find good resources that can be used in a collection to help understand a complex topic. While analyses of how people use physical books has suggested that multiple views of texts could be useful for complex sensemaking work [14, 1], few studies of how such an interface could be built and used have been conducted since then. Our work fills this gap in the research.

Theories of Learning

Marton's Variation Theory, as summarized by Suhonen et al. [18], is defined by the dimensions of variation necessary to fully communicate a concept to a student: *contrast* ("in order to experience something, a person must experience something else to compare it with"); *generalization*, or the ways something can vary without becoming something else; *separation*, or looking at the variation only across specific features; and *fusion*, where multiple critical aspects of the concept are varied simultaneously. In other words, variation reveals which aspects of a phenomenon are superficial/irrelevant and which are innate/critical to its definition [9]. It is a framework that now guides the design of some critical reading exercises [21] and exercises for novice programmers [5].

User Interface Design

Tufte pioneered a layout technique called "small multiples," designed to help viewers make rapid decisions about a wide array of items or variables: "Small multiple designs ... answer directly by visually enforcing ... the differences among objects, ... the scope of alternatives." We took inspiration from both of these layouts when designing the flowing grid layout for DocMatrix.

The common term sidebar was inspired by Hearst's faceted browsing [25]. But rather than display facets derived from metadata about each document, we extracted the common terms from a source closer to the content of the documents themselves: the tables of contents. Clicking on any of the terms in this list exposes the tables of contents, with relevant chapter titles highlighted; the actual terms contained in the tables of contents are the key to this ontological alignment.

Grokker is a document-clustering visualization system, with small popup windows to read texts in parallel [17]. Unlike DocMatrix, Grokker's primary representation of a corpus of documents is as clusters of dots, but the study design and results are still relevant here. The task for Grokker readers was to quickly browse a large document collection, and then answer a set of questions to test their understanding. A key finding of this study was that small details of document viewability and the amount of time it took the participants to access content dramatically affected how much they understood about the domain. In other words, small changes in the amount of time to switch between related documents was an important variable.

[3] describes a case study of the Reflective Online Searching System (ROSS), designed to help students learn how to search and develop meta-cognition about searching. ROSS allows students to search online and reflect on the variation in information they come across. However, their interface does not appear to explicitly align resources and put them side by side, for easier skimming and comparison, like DocMatrix.

DESIGN GOALS

When a user is trying to teach themselves something using results from an online search or a search through a library's collection, they face several challenges. First and foremost, they must find information that is at the appropriate level, given their prior knowledge. For difficult topic areas like machine learning, users may need to find multiple explanations and figures, from different sources, in order to fully understand an idea, as described by [11].

We hypothesize that readers will get the most benefit from features that facilitate switching between and comparison across texts, which enable both (1) finding the subset of resources that are at a level appropriate for the user and (2) finding multiple explanations for the same idea. We also hypothesize that readers will find more helpful information with DocMatrix, relative to a traditional online book browsing experience like Google Books.

Our design goals are to explore: (1) features that vastly improve a reader's ability to explore and exploit multiple books and (2) how a tool with these features would change the reader's experience of teaching themselves.

THE DOCMATRIX UI

To fulfill our design goals, we built the DocMatrix interface for Google Books. DocMatrix is a UI for a collection of structured documents that supports simultaneous multiple-document reading. In the context of Google Books, it's not simply a collection of individual e-book reader panels; it also has an analysis system and search UI that spans the entire collection of books. The grid interface is specifically intended to promote multi-reading, rather than uni-reading, as a way to enhance the reader's search and discovery process. While it allows for focused reading of one book, it dramatically lowered the cost of switching to reading another book's treatment of the same idea [15]. When readers switch across books easily, they may be more able to find information from different points of view and at an appropriate level of sophistication, given the reader's prior knowledge. DocMatrix for Google Books is comprised of: (1) a grid of embedded books, (2) a sidebar of common terms derived from all the books' tables of contents, and (3) tables of contents for each book with clickable chapter titles that open the book to that chapter and expose snippets from that chapter containing the query or term of interest.

1. A Grid of Embedded Books Readers submit a query at the top of the screen, as they might for any search engine. As shown in Figure 1, the top forty most relevant books, as ranked by Google, are embedded into the page, each in their own viewer, in a three-column wide grid, offering a "small multiples" style of display [22]. Each viewer pane allows readers to scroll through, zoom in and out, and also perform full-text searches on all the viewable pages of the book.

In the system's initial state, two rows of books are within view, with the rest of the collection just below, off the screen but reachable by scrolling. The embedded books are readable in place, without zooming in or switching contexts to a new window, on both a 15" retina screen and a 30" LCD screen. If the user wants to focus in on only one book, then they can use either the viewers or the browser's zoom function.

The book viewers' size supports glanceability in addition to readability. Viewers are initialized to show each book's cover. The user's first impression is an array of large book covers, with metadata, such as title, authors, publication date, and general description or relevant snippet, shown beneath. Many book covers are explicitly designed to communicate the author's approach, perspective, or intended audience. The cover design may also suggest a time period of publication. For example, when a user searches for books on algebra, the first edition of Euler and Lagrange's *Elements of Algebra* shows up as a smudgy scanned title page alongside a plain, clearly dated, mono-chrome cover of a graduate algebra textbook from 1974. Readers can easily pick up clues, at a glance, about what the book offers.

2. Common Term Sidebar Beside the grid of embedded books, DocMatrix displays a sidebar of common terms, ordered by frequency. To construct this sidebar, each book's table of contents is treated as a bag of tokens, stripped of stopwords, and then combined with the token bags of all the other tables of contents in the collection. The counts of each term in the resulting bag determine their order in the sidebar and are displayed in the sidebar next to their associated term.

Clicking a term in the sidebar (1) filters the collection down to just books with that term within their chapter titles, (2) appends tables of contents below each book, and (3) highlights all occurrences of the term within both the chapter titles and the full text of each book. The resulting filtered set of books are still shown in a 3-column layout, but once their tables of contents are appended, the blocks of content representing each book have potentially different heights. Rather than fitting into a grid, they are allowed to take up as much vertical space as they need and are laid out in a "Pinterest" style.

3. Enhanced Tables of Contents The tables of contents beneath each book behave like accordion menus, with chapter titles as the top-level menu items. Clicking on a chapter title (1) opens the corresponding book to the start of that chapter and (2) exposes a submenu containing extracted snippets from the chapter. A tally next to the chapter title indicates how many snippets are in that submenu. This is illustrated in Figure 2. If no sidebar term is selected, the snippets are all mentions of the original query within the chapter. If a sidebar

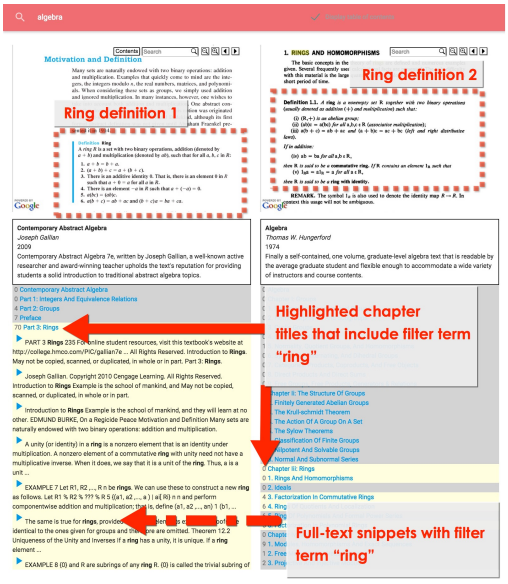

Figure 2. Snapshot from a reader's in-depth investigation of rings in abstract algebra, using DocMatrix.

term is selected, the snippets are all mentions of the sidebar term within the chapter. Clicking on a snippet also opens the book to the page containing the snippet. Finding the snippet within its page is relatively easy because all occurrences of the sidebar term are highlighted.

The snippets and snippet counters for each chapter give readers some X-ray vision into a chapter. Consider a reader who clicks on the term *rings* within the common term sidebar. The books are filtered down to just those books with chapter titles that include the term *rings*. These filtered books' tables of contents are exposed, and chapter titles mentioning *rings* are highlighted. There may be additional chapters that mention *rings* relatively frequently, but do not happen to mention it in the chapter title. Including the counts of snippets mentioning *rings* within each chapter allows the reader to see which chapters actually discuss rings a great deal, regardless of title.

Alternative Designs

Between pilot sessions, we iteratively added and removed features to the grid of embedded books. These features included the drag-and-drop reordering of books, liking or starring a book, hiding or removing a book, sorting books by metadata, methods for curating a working set of books, and popping books out into separate tabs. However, these features only distracted readers from the intended behavior of reading about a particular topic across multiple documents. The only feature that supported this behavior explicitly was the sidebar of common terms, which remains.

Example Domains

In Figures 1 and 2, we see snapshots of a reader interacting with DocMatrix, using the query *algebra*. Figure 1 shows the initial gestalt view of books, with prominent covers, and Figure 2 shows a clip of the DocMatrix UI after the user has delved deeper into a subtopic, rings. Figure 2 specifically

shows two books embedded side-by-side in the interface after the sidebar term *rings* has been clicked. Their tables of contents are displayed. Chapters with *rings* in their titles are highlighted. The reader has clicked on the first highlighted chapter in both books, exposing any extracted snippets on *rings* within those chapters in plain text and opening both books to the clicked chapters. Each author's definition of *rings* is in view, for easy comparison. The reader may find one definition clearer than another, given their prior knowledge, or feel that the cumulative effect of the two definitions and accompanying explanations clarifies the idea of rings more than any one explanation could do alone.

The query *cancer* changes the domain of books and the user's experience. The collection of 40 top books returned to DocMatrix by the Google Books API includes a comprehensive guide to cancer treatment by the Mayo Clinic, an exposé on the 'war on cancer,' a doctor's guide to avoiding cancer in the first place, a well-known literary work centered around cancer, a textbook on cancer biology, and an astrology book that mentions cancer as an astrological sign. The common term sidebar reveals that the terms *breast*, *treatment*, and *diagnosis* are occur frequently in these books' tables of contents. By clicking on *treatment* in the sidebar, the books are filtered down to just those containing chapters with 'treatment' in the title, such as 'Understand Your Conventional Treatment Options,' 'Treatment Side Effects,' and 'The Rational Treatment of Cancer.' With their tables of contents expanded and shared term highlighted, the reader can focus on what facets of cancer research they wish to pursue, while simultaneously keeping books with other foci just a click away.

STUDY DESIGN

Prompt Design

In our screener for participants, we screened for: (1) nearly uniform unfamiliarity across all of the participants, (2) consistent terminology across books, (3) depth or richness and (4) not completely out of the intellectual reach of participants. Unfamiliarity is necessary because we want participants to teach themselves something they have not been exposed to before. Consistent terminology allows subjects to more easily look at explanations of the same concept across books without having to build synonym models or entity resolution into the prototype. Subject depth is necessary so that participants are engaged in their exploration of the subject for the entirety of their session. While a more obscure subject in history might satisfy (1), (3) and (4), it is unlikely to satisfy (2)–consistent terminology across books. For that reason, we only considered topics in STEM fields–science, technology, engineering, and mathematics.

Abstract algebra fit all three of these criteria: (1) Few subjects were familiar with any concepts from abstract algebra. (2) The concepts have specific technical names that are consistent across all books on the subject. (3) It has enough depth for an entire course to be taught on the subject. (4) It is not necessarily out of the intellectual reach of participants: while it is usually taught in college-level mathematics courses, it has also been formally taught in high school mathematics classes by teachers like Benjamin Sapolsky, as described by one of his students, Martin Glassman (personal communication). We believe, therefore, that it is approachable by non-math majors in our pool of subjects, if they can use the interfaces given to them to find the right books to read.

Prompt We designed the following prompt: "Explain something about rings, groups, and fields in algebra." It references three major concepts from abstract algebra. If one concept is particularly confusing to the subject, they can move on to another.

Procedure

Each subject participated in a 60-minute session that began with a short, scripted 10-minute tutorial on how to use both the DocMatrix and standard Google Books interfaces. This prepared them for the 30-minute learning activity that followed, where the participants would use one interface for 15-minutes, then switch to the other interface style and continue learning. The order of interfaces within the learning activity was counterbalanced across subjects. Using both interfaces for the same learning activity enabled subjects to make direct comparisons in the (approximately 10-minute) post-study reflection survey that immediately followed.

During the session and post-study survey, DocMatrix was referred to as "Book-Grid" and Google Books was referred to as "Book-List." These terms were memorable and called out the main visual differences, without biasing subjects by implying that one was specifically engineered for helping readers learn.

For the learning activity, subjects were given the prompt described in the previous section. We were effectively asking them to teach themselves as much as they could about three core concepts in abstract algebra in half an hour. To control for subjects' individual abilities to refine queries, we asked that they only retrieve book results using the query *algebra*. They were not allowed to consult other resources.

Subjects were asked to collect cropped screenshots of portions of book pages that gave them an insight or clarified a point (e.g., a formal definition of an idea, an illustrative example, a context-giving short passage, or a chart) during the learning activity. Subjects were trained in how to use a corporate Chrome extension for this purpose during the tutorial. Subjects saved cropped snapshots in nearly constant time that was consistent across subjects. In descriptions that follow, we will refer to this as 'saving' X, where X was the captured portion of a page. An additional, unintended benefit of this method is that it allows subjects to decisively indicate what they found helpful, rather than putting the burden on the experimenter to infer an insight from verbal expressions like "Oh . . . !"

Subjects were alternatively assigned to one starting interface or another in the order that they arrived in the lab. Of the nine subjects, four used the DocMatrix interface first and five experienced the Google Books interface first, switching to the other interface at the halfway point of the learning session. The session was conducted on a 30" screen (2560 x 1600 pixels), driven by a MacBook Pro, which maximized the size of embedded books.

Control Interface

We chose Google Books as the control interface. The books displayed in Google Books are *the same books in the same order* as those displayed in DocMatrix. In the Google Books interface, books are represented by a small thumbnail of the book's cover alongside its title, author, publication year, and general description. Clicking on a title switches the user into a new context: a viewer for that book alone. The viewer shows a long scrollable set of pages. Unviewable pages are collapsed into a thin bar that indicates the pages are outside the publisher's preview. The reader's original query is highlighted in the viewer. In a thin banner, readers can click on links to step through all the query occurrences within the full text. They can also switch the view to a list of book page snippets, ordered by either relevance or page number. The user can search the book for new terms and phrases or click on "About This Book" to see book metadata, related books, selected page snapshots, the table of contents, common terms and phrases, and popular passages.

Measurements

We measured seven kinds of behavioral data from each subject's session. (1) *Number of snapshots made*: we opened the notes of each subject for each interface and counted the number of book snapshots were saved there. As per the instructions given, each saved snapshot represented a helpful book portion that the student found in that book. (2) *Verbal discovery comments*: we recorded the subjects talking aloud to the experimenter while interacting with each interface. Comments ranged from simply reading book paragraphs aloud to verbal expressions indicating discovery, e.g., "Ah ha!," to complaints and praises for specific interface features. (3) *Books that came into view*, (4) *how many books were interacted with*, (5) *how many chapters the participant read from*, (6) *how many full-text searches performed*, and (7) *how often any term was clicked in the common term sidebar*. (For the measures 3-7 we recorded the screen and then manually counted each behavior from the screen recording.) These counts and transcriptions did not demand the kinds of judgement calls that require multiple raters for confirmation. Lastly, we analyzed subjects' responses in a post-study reflection survey, which included both free response and Likert scale questions.

RESULTS

Participants

We recruited nine subjects (5 women and 4 men), with a mean age of 20.8 (σ = 1.3), from a large intern mailing list at a major tech company. They collectively covered a range of mathematical comfort and knowledge. The majority of our subjects were comfortable with the topic, having taken calculus, linear algebra, and perhaps one or two additional classes on statistics, discrete math, or differential equations. While the sample size of 9 participants is small, our goal was to understand reader behavior with this type of interface by watching their behaviors and doing a qualitative analysis of their performance. Since we are looking for behaviors that would hold broadly, we felt the small sample size was justified [23].

Quantitative Results Regardless of which interface they used first, every subject saved just as many or more (50% on average) helpful book passages with DocMatrix as with Google Books. This difference is significant by a paired t-test (p = 0.016). The average number of saved helpful passages with DocMatrix was 4.7 (σ=2.0) and, with Google Books, only 3.2 (σ=1.3).

Two-thirds of the subjects (6 of 9) interacted with strictly more books in the DocMatrix condition than those in the Google Books list condition. The remaining three subjects were exceptions, each in their own way. S5 did not understand what the DocMatrix interface was doing and why, as revealed in her post study reflections. She interacted with fewer books in DocMatrix than Google Books. S9 used Google Books' features for discovering more relevant books better than any other subject, i.e., by opening an individual book's 'About' pages, then perusing the Related Books section. She was also the only subject to look beyond the first page of ten book results. She interacted with more books in the Google Books condition. Finally, S6 was a math major who had already taken a class in abstract algebra. Regardless of what books he opened, he could follow most of what he read. He interacted with the same number of books in both of the interfaces. Since S6 was sophisticated in the subject area, good book choice mattered less for S6 than for people with less background in the topic.

Both interfaces afford searching for terms within a particular book. DocMatrix, however, also allows readers to filter and highlight a collection of books with a term from the common term sidebar. When limited to only searching within a particular book (Google Books condition), subjects searched an average of 7.7 (σ=3.8) times. In the DocMatrix condition, the number of single-book searches dropped significantly ($p < 0.001$, paired t-test) to 1.3 (σ=1.7). Instead, readers took advantage of the cross-book filter and highlight mechanism, clicking terms in the common term sidebar an average of 3.4 (σ=2.1) times in their 15-minute session with the interface. These averages indicate that when subjects [using DocMatrix] had the ability to simultaneously filter and search books for a term as well, they used it in addition to or instead of full-text search within a particular book.

Orienteering Behavior In each interface, subjects needed to hunt for information they could make sense of, and did so often by moving from topic to topic across texts in an orienteering style [13]. S1 literally called it a "treasure hunt." Regardless of interface, every subject collected definitions for one or more of the three terms in the prompt: ring, group, and field. S9 started her learning session by announcing: "The first thing I want to do is figure out the definitions of these terms." However, how subjects went about collecting their insights from these books was significantly affected by the interface they were using.

In the Google Books interface, all but one subject stayed on the first page of search results. This first page of search results shows the first ten of the top forty books retrieved for the same query in DocMatrix. Some subjects were biased toward the top results in that list. In this way, the spatial distribution of

books in the results list had a significant effect, favoring those at the top of the list.

However, in DocMatrix, subjects clicked on sidebar terms in order to filter those top 40 books down to the subset that had chapter titles mentioning the selected term. Since that subset of books was displayed with their tables of contents highlighted to draw attention to all the chapters mentioning the term, subjects, as a general rule, steered toward opening books to the *first* highlighted chapter in each book they interacted with.

S8 talked aloud about selecting books to interact with based on where the highlighted chapters occurred within the tables of contents of each book. S8 was betting that highlighted chapters early within any book would be more approachable than those occurring later in any other book.

Both S1 and S8 spoke about using textual cues to select chapters to open in DocMatrix. S8 clicked on the sidebar term `groups`, scanned the top 5 books in the resulting filtered book set, and quickly saw a promising chapter title in the fifth book, exclaiming, "'Definitions and Examples of Groups'—This looks great!" S1 responded to the chapter title "Groups, first encounter" with the phrase, "You seem friendly!" She clicked on the chapter to open the book to the chapter's beginning and simultaneously expose snippets from the chapter's full text to determine whether it was actually going to help her.

Within the Google Books interface, the differences between algebra textbooks were not as salient, textually. S8 commented, "Most of the books are pretty much just called algebra which makes it hard to know which would be the most useful. I'm . . . seeing if they say 'introduction,' 'first course,' [or] 'for undergrads.'" S7 opened a book "because it says 'a concrete approach' in the title. Maybe it will have examples." S6 looked at books' descriptions, noting that while most were for graduate students. He remarked that "an undergrad text would be better, because you're assumed to know nothing."

There were two visual aspects of the books that subjects attempted to use while orienteering their way to the information they found helpful: book covers and book pages. In the Google Books interface, books have only a small thumbnail of a cover. Both S1 and S8 stated that they gleaned little information from this; the thumbnails were too visually similar to help with the decision of which book to open next.

When S1 clicked on a particular book in Google Books, it loaded in a new window, only to reveal that the book was old enough to be a scanned copy rather than a publisher's digital version. The text was smudged. S1 immediately closed it, dismissing it by saying, "This looks old. I don't want this. [laughs] It looks hard to read!" A different book in the displayed results was formatted well; it drew readers' attention to important points by, for example, putting definitions in colorful boxes. S8 expressed a preference for more recent books, assuming that newer books are more readable through better formatting, book image quality, and/or writing style. In DocMatrix, books' covers are much larger to begin with, and readers can scroll within any book to see the formatting and image quality without leaving the context of their book collection. S5 specifically praised these large covers as one of DocMatrix's strengths.

In addition to visual, textual, and spatial cues, subjects used search tools to identify good books. The browser-based within-page search feature that subjects invoked using the shortcut `Ctrl+F` was employed by subjects in both interface conditions. S1 tried to search the first page of book results this way, looking first for rings, then fields. Since no book's metadata displayed in the Google Books results page specifically mentioned either of these terms, there were no hits. It was not, as S5 requested aloud, searching all the books' full text for hits on a term of interest. In DocMatrix, S1's `Ctrl+F` for a term of interest searched all the books' tables of contents and snippets of expanded chapters as well. DocMatrix's behavior is closer to what S5 explicitly asked for than Google Books' and more likely to return a result.

Text densely populated with technical terms was impenetrable to those not already familiar with the terminology. In Google Books, most subjects reacted to impenetrable text by abandoning the book and moving along to what appears to be the next best option, based on metadata visible in the search results page. In DocMatrix, subjects used the common term sidebar to quickly find and switch to another book chapter with more appropriate explanations, given the subject's mathematical background. For example, S8 encountered a passage that was riddled with unfamiliar jargon. She said aloud, "I'm going to look at some of the other books and see if they explain it in less mathy ways."

When subjects found a book they believed to be helpful, some responded emphatically. S1 clicked on three sidebar terms in rapid succession in the DocMatrix interface. After each click, a different subset of books would appear, but one book stayed in view the whole time. She realized that there must be chapters on all three terms in that book, exclaiming, "Groups, rings, fields, they're all here . . . This book is love."

Finding Relationships Between Concepts Many subjects sought to go beyond definition gathering to grasping the relationships between rings, groups, and fields. At the end of the half-hour learning session, S7 clicked open all the chapters on fields he had recently looked at across several books, going back and forth between them, obviously searching for something. When prompted to explain, he said, "I was looking for a definition that would . . . tie together how rings, groups, and fields were related." While also using DocMatrix, S8 found the introductory paragraphs of a particular book's chapter on groups to be worthy of saving, just for the context it provided: "It talks about how [a group] relates to other things I'm learning about. I think it has the broad knowledge that I might be able to look at this and look things up." S1 saved the sentence "To this day, matrix groups over finite fields are among the most important classes of groups" because "it gives me a lot of context as to the relationships between these concepts." One book explained the relationship between rings and fields with a chart: rings were on one end and fields were at the other, connected through intermediate boxes filled with more

specialized terms. Several subjects saved it, despite not understanding all the details.

Subjects followed several different strategies for finding descriptions of these relationships in understandable terms. For example, S4 appreciated the counts next to each term in the common term sidebar; she felt that it suggested relationships across terms with similar counts: "Knowing that group and rings were similarly common made me think they might be related." She also tried clicking the sidebar term *groups*, then using the within-book search function to look for 'fields.' While DocMatrix made it easy to search across books for a single term, we did not observe any significant difference between DocMatrix and the Google Books interface's support for finding relationships between terms.

Switching Between Books Subjects appreciated the ease of switching between books in DocMatrix. While using the Google Books interface, S1 said, "It makes me crave [DocMatrix]! . . . because I keep wanting to jump between books but it's not very easy to do that [in the Google Books interface]." In the post-study reflection survey, S7 praised DocMatrix for "the ability to quickly switch between books without having to navigate through a series of pages."

Comparing Across Books With a collection of books at their fingertips, each with their own unique presentation of the key ideas in abstract algebra, eight out of nine subjects saved definitions for rings, groups, and fields *from multiple sources*. S9 explained that "one of the things I like doing is just figuring out different ways people will explain something." While using DocMatrix, S1, S8, and S9 explicitly weighed two definitions of a concept from two different books and, given some internal set of their prior knowledge and preferences, decided which one was better for them to save. S1 imagined going further, if she had more time: "Using this technique [of collecting snapshots], which I might start adopting, I'd probably dump all the definitions of rings that I find useful, from multiple books onto one slide just so I can compare them. I have two definitions right now, one which doesn't make a lot of sense to me and this one which just seems to make a little more [sense] . . . Having them all in one place would probably make it a little easier."

Interface Features The most prominent feature that differentiates DocMatrix from Google Books is the grid of books whose full text is available within the page, just by scrolling inside each book viewer. Five out of nine subjects praised DocMatrix's ability to show multiple books at once. S1 wrote, "Having everything all in one place was amazing. Opening extra windows and switching between tabs is one of the most annoying things about doing web searches for specific queries." S9 observed, "I liked that I could see all of the different books and their contents without having to explicitly click on one book. I think this made it easier to scan books to see if they were interesting enough to further look for more information."

There were some reservations about this feature, as well. S6 explicitly wrote that they did not appreciate how each book they focused on was "still part of the overall landscape." S8 also pointed out that having three columns of book viewers, even on a large monitor, constrained the books themselves to be "kind of small." She felt it necessary to zoom in and out frequently as they switched between books, and S9 did not find the browser-based zoom interaction intuitive. S6 felt "forced" to view multiple books simultaneously; he did not learn or remember the browser-based method of focusing on a single book within DocMatrix. S8's praise for side-by-side book viewing was tempered by feeling "a bit overwhelm[ed]" by the volume of information in one page.

In the traditional Google Books interface, some subjects found the simplicity of a single scrollable book in view both a blessing and a curse. S4, S6, and S8 all expressed an appreciation for seeing multiple pages of a single book in view at once on the large 30" screen. It helped S7 focus. S5 and S9 appreciated that the interface is externally consistent with other online searching and reading experiences. The downside of this mono-book reading environment was the lack of support for comparing content across books, which S1, S2, S4, S5, and S9 all observed. S4: "[It was] hard to tell if books were covering the same material or not. I had to open each one to check."

The continuous visibility of the table of contents for each book supported reader's orienteering. S4 explained how this prominence set DocMatrix apart from Google Books: "I was quickly able to see which chapters had those terms, and even how far up in the table of contents it was. If 'field' was in the first half of the table of contents, I thought it might be an introductory chapter that would be helpful to check out."

According to the talk-alouds of several subjects, the common term sidebar fulfilled one of its intended purposes: making it easy and quick to simultaneously access information from multiple documents about a named concept. After first using the Google Books interface, S4's first click in the DocMatrix interface was on `groups` in the common term sidebar. S4's second click was on the first highlighted chapter of the first book in the new filtered set. With just those two clicks, she was reading an introductory paragraph on rings. After saving that helpful paragraph, she said, "That was a lot easier . . . just finding information about groups. It was a lot quicker than I thought it would be." S2, S3, and S7 all praised how they could "quickly jump to different topics" by clicking on sidebar terms. S1 clicked on `rings` in the sidebar, and when all the books with all the chapters on rings were highlighted, she saw that a book's content was more relevant than she had guessed: "Wow! Wow . . . [this book has] a lot more about rings [than I expected] . . . I should have bet on you." S2 believed that "being able to sift through multiple books of data at once significantly improved research efficiency." The common term sidebar was such a prominent feature that he forgot about the small search box at the top of each book viewer.

The second intended purpose of the sidebar was to give a composite table of contents for books retrieved by the reader's query. S1 explicitly mentioned this value of the sidebar in their post-study reflection: "At a glance, you could tell what books on the subject would be focusing on and what areas you want to consider looking into first."

Some conceptualized the sidebar as pre-processed query suggestions, rather than a reflection of the terms common across all books. Specifically, S6 called the terms in the common term sidebar "sample queries," and S4 remarked that DocMatrix made it "very easy to search for terms."

In the post-study survey, six out of the nine subjects opted to use DocMatrix instead of Google Books to teach themselves something in the future. Subjects also responded to a battery of Likert scale questions on a 7-point scale (1-strongly disagree, 7-strongly agree):
1. Subjects most strongly felt that DocMatrix "helped me bring together information from multiple sources and points of view" (5.7, σ=1.2) more than Google Books (3.4, σ=1.3).
2. Subjects also strongly felt that DocMatrix made it "easy for me to tell whether a book was useful or not" (4.8, σ=1.6) compared to Google Books (2.7, σ=1.5).
3. There was a moderate difference in subjects' perception of their ability to find useful information with each interface (DocMatrix: 5.7, σ=0.9; Google Books: 4.4, σ=1.0).
4. There was no difference in how subjects perceived their ability to "find information at the appropriate level, given my prior knowledge" or "learn what I set out to learn."

DISCUSSION

DocMatrix is designed for readers to both explore a collection of documents on a topic and exploit those most appropriate, given the user's goals and prior knowledge. The prototype we studied, DocMatrix for Google Books, fulfilled this goal in multiple ways: (1) providing the reader with perspective about the common and uncommon terms in all the tables of contents, (2) helping readers pick which books to actually read by making the books' full text and tables of contents available directly on the page, and (3) reducing the cost of reading about concepts and terms across multiple books. The author-generated table of contents, i.e., a structured annotation of the content, was the key to making this method work for the books corpus.

The average subject using DocMatrix interacted with more books and, as a result, found books that were more helpful to them. This is evidenced by the increased number of helpful book passages that subjects saved for their notes in the DocMatrix condition. Using DocMatrix, subjects jumped right into considering which chapters from which books to read. In the traditional Google Books condition, subjects had limited information with which to pick a book to read, and then it is more difficult to switch between books once that decision is made. This is suppored by participants' subjective evaluations: using DocMatrix, subjects felt they could more easily assess the books' usefulness and synthesize knowledge across books.

DocMatrix reduced the cost of switching between texts, so users could more easily handle reading situations where the text became impenetrable or unviewable. If a reader encountered a passage they could not understand or a publisher-hidden page, they could open another book to the same topic in one click. In other words, based on their behavioral data, participants had less tunnel vision for only one resource when using DocMatrix.

The common term sidebar in DocMatrix offered a qualitatively different kind of experience, i.e., search across a selected collection. Rather than search within books one by one, subjects clicked on terms in the sidebar to filter and highlight all the books in the collection. This was followed, on occasion, by traditional full-text searches within particularly promising books.

Limitations

The common term sidebar currently only includes 1-grams extracted from the documents' tables of contents. This helped readers find rings, groups, and fields in abstract algebra books, but might break down when a reader wants to brush up on the "Battle of the Bulge." Future DocMatrix implementations can use more sophisticated entity retrieval techniques to populate the sidebar.

The sidebar's effectiveness also depends on the quality of the tables of contents. Some books have very generic chapter titles, e.g., "Introduction" or "Related Work," and some have creative chapter titles that do not specifically mention key ideas within them. As long as there are enough books, the sidebar will reflect common ideas across books. However, after the term in the sidebar is clicked, books with chapters on a topic that do not mention it specifically in a chapter title will get filtered out of view. This did not affect any of our study subjects, but, in the future, one could imagine filtering and highlighting books based on both chapter titles and chapter content.

Future Work

Future work includes adapting DocMatrix for more sources of structured documents, including Wikipedia, newspapers, and captioned videos. As long as a document or video has (1) section labels akin to a table of contents and (2) a viewer or player that can be started at arbitrary points, the DocMatrix design can be applied. Documents without structure could be pre-processed to extract section titles.

The common term sidebar could also become more of a true composite table of contents. Rather than ordering terms by counts, they could reflect the average ordering of terms within the tables of contents in the book collection. For example, perhaps, in 95% of the books in the algebra collection, groups appear before rings and rings appear before fields in the tables of contents. Order matters to learners, since it may reflect increasing complexity, prerequisites, and dependencies.

DocMatrix does not yet address the cognitive challenges of jumping into an arbitrary section of a book. Future versions of DocMatrix could help readers trace back through a whole thread of terms and explanations mentioned in a document, so they could read only the slice of the document that they need to read.

Conclusions

DocMatrix is an interface for viewing multiple sources in parallel, with three key features: a grid of viewers, a common term sidebar, and enhanced tables of contents. It enables an interested learner to view, filter, highlight, and search many

documents on a topic simultaneously. When using the DocMatrix interface for Google Books instead of a traditional single-book interface, readers could find, read, and synthesize more information across multiple books.

REFERENCES

1. Annette Adler, Anuj Gujar, Beverly L Harrison, Kenton O'Hara, and Abigail Sellen. 1998. A diary study of work-related reading: design implications for digital reading devices. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM Press/Addison-Wesley Publishing Co., 241–248.
2. Benjamin S Bloom. 1984. The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational researcher* (1984), 4–16.
3. Christine Bruce, Sylvia Edwards, and Mandy Lupton. 2006. Six Frames for Information literacy Education: a conceptual framework for interpreting the relationships between theory and practice. *Innovation in Teaching and Learning in Information and Computer Sciences* 5, 1 (2006), 1–18.
4. Richard Cox and Paul Brna. 1995. Supporting the use of external representations in problem solving: The need for flexible learning environments. *Journal of Artificial Intelligence in Education* 6 (1995), 239–302.
5. Anna Eckerdal and Michael Thuné. 2005. Novice Java programmers' conceptions of object and class, and variation theory. In *ACM SIGCSE Bulletin*, Vol. 37. ACM, 89–93.
6. Susan R Goldman. 2004. Cognitive aspects of constructing meaning through and across multiple texts. *Uses of intertextuality in classroom and educational research* 317352 (2004).
7. James G Jardine. 2014. Automatically generating reading lists. *University of Cambridge, Computer Laboratory, Technical Report* UCAM-CL-TR-848 (2014).
8. Harvard Library Innovation Lab. 2012. Stack View. (2012).
9. A Leung. 2012. Variation and mathematics pedagogy. In *Proceedings of the 35th Annual Conference of the Mathemafics Educafion Research Group of Australasia*. 433–440.
10. Kurt Luther, Nathan Hahn, Steven P Dow, and Aniket Kittur. 2015. Crowdlines: Supporting Synthesis of Diverse Information Sources through Crowdsourced Outlines. (2015).
11. Metacademy.com. 2012. How to learn on your own. (2012).
12. Aditi Muralidharan, Marti A Hearst, and Christopher Fan. 2013. WordSeer: a knowledge synthesis environment for textual data. In *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*. ACM, 2533–2536.
13. Vicki L O'Day and Robin Jeffries. 1993. Orienteering in an information landscape: how information seekers get from here to there. In *Proceedings of the INTERACT'93 and CHI'93 conference on Human factors in computing systems*. ACM, 438–445.
14. Kenton O'hara and Abigail Sellen. 1997. A comparison of reading paper and on-line documents. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*. ACM, 335–342.
15. Daniel M Russell, Malcolm Slaney, Yan Qu, and Mave Houston. 2005. A Cost Structure Analysis of Manual and Computer-supported Sensemaking Behavior. *Proceedings of Intelligence Analysis* 2005 (2005).
16. Dafna Shahaf, Jaewon Yang, Caroline Suen, Jeff Jacobs, Heidi Wang, and Jure Leskovec. 2013. Information cartography: creating zoomable, large-scale maps of information. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1097–1105.
17. Malcolm Slaney and Daniel M Russell. 2005. Measuring information understanding in large document collections. In *System Sciences, 2005. HICSS'05. Proceedings of the 38th Annual Hawaii International Conference on*. IEEE, 105–105.
18. Jarkko Suhonen, Janet Davies, Errol Thompson, and others. 2007. Applications of variation theory in computing education. In *Proceedings of the Seventh Baltic Sea Conference on Computing Education Research-Volume 88*. Australian Computer Society, Inc., 217–220.
19. HJM Tabachneck, KR Koedinger, and MJ Nathan. 1994. Toward a theoretical account of strategy use and sense-making in mathematics problem solving. (1994).
20. Alice Thudt, Uta Hinrichs, and Sheelagh Carpendale. 2012. The bohemian bookshelf: supporting serendipitous book discoveries through information visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1461–1470.
21. Siu Yin Annie Tong. 2012. Applying the Theory of Variation in Teaching Reading. *Australian Journal of Teacher Education* 37, 10 (2012), 1.
22. E.R. Tufte. 1990. *Envisioning information*. Number v. 914 in Envisioning Information. Graphics Press.
23. Carl W Turner, James R Lewis, and Jakob Nielsen. 2006. Determining usability test sample size. *International encyclopedia of ergonomics and human factors* 3, 2 (2006), 3084–3088.
24. Sam Wineburg. 1998. Reading Abraham Lincoln: An expert/expert study in the interpretation of historical texts. *Cognitive Science* 22, 3 (1998), 319–346.
25. Ka-Ping Yee, Kirsten Swearingen, Kevin Li, and Marti Hearst. 2003. Faceted metadata for image search and browsing. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 401–408.