To Print or Not to Print: Hybrid Learning with METIS Learning Platform

Joshua Hailpern[1], Rares Vernica[2], Molly Bullock[3], Udi Chatow[2], Jian Fan[2], Georgia Koutrika[2], Jerry Liu[2], Lei Liu[2], Steven Simske[4], Shanchan Wu[2]

[1]HP ArcSight & HP Labs, [2]HP Labs — *Palo Alto*, [3]HP, [4]HP Labs — *Fort Collins*

[1] Sunnyvale, CA, [2] Palo Alto, CA, [3] Boise, ID, [4]Fort Collins, CO

{joshua.hailpern,rares.vernica,molly.bullock,udi.chatow,jian.fan,koutrika,jerry.liu,lei.liu2,steven.simske,shanchan.wu}@hp.com

ABSTRACT

As part of the explosion in educational software, online tools, and open educational resources there has been a rapid devaluation of printed textbooks. While digital texts have advantages, printed textbooks still provide irreplaceable value over online media. Therefore technology should enhance, rather than eliminate printed text. To this end, this paper presents METIS, a hybrid learning software/service platform that is designed to support active reading. METIS provides easy digital-to-print-to-digital usage, simple creation of Cheat Sheets & FlexNotes for personal note taking and organization, and a custom flexible rendering & publishing engine for education called Aero. METIS was designed based on lessons learned from a formative study of 523 students at SJSU, and validated through focus groups involving 32 educators and students at both high school and college levels.

Author Keywords

Education; Active Reading; Print; Hybrid Learning

ACM Classification Keywords

K.3.1 Computer Uses in Education

INTRODUCTION

The ubiquity of computers, portable electronics, and Internet access has spurred changes in education. Billions of dollars are being invested by educational institutions and companies. There is a rapid increase in online learning tools, social networks integration and mobile device support.

In parallel, there has been an explosion in the learning management platform and open educational resources spaces. In this trend of "digitizing" education, there has been a rapid devaluation of printed textbooks. However, printed textbooks still provide irreplaceable value. Instead of eliminating this value, education should be based on print *enhanced by* the use of technology, rather than replaced by it. Some of the

Figure 1: Hybrid Learning Model
Digital to Print to Digital

advantages of printed textbooks include: 10-30% faster reading rate [24], lack of distractions, no device compatibility or internet connection issues, more cost effective, and, most importantly, are the medium still preferred by the majority of the students [18]. On the other hand, digital content allows for more interactive and collaborative learning, robust search, and the inclusion of multimedia content.

This re-examination of printed and digital education texts has given rise to a set of interaction techniques called *Active Reading* [23]. When students engage in Active Reading, they highlight, annotate, outline, search content and compare pages. Unlike passive "pleasure" reading, this Active Reading process has been shown to be critical to the education and learning process. Most importantly, education systems must support this class of interactions.

In this context, there is a surprising lack of integrated printed reading material by many research systems (e.g. XLibris[22], and LiquidText[34]), and when printing is included, proprietary hardware/pens & paper is required, limiting the use and increasing costs (e.g. Papiercraft [14]). In response, this work focuses on engineering a seamless[1] interaction with paper and digital content, and the transition between the two. Specifically, we created **METIS** a hybrid learning software/service platform that is designed to support active reading. The primary contributions of METIS are **1) Hybrid Learning** platform for digital-to-print-to-digital, **2) Cheat Sheets & FlexNotes,** a personal note taking and organization system facilitating the hybrid learning model, and **3) AERO**, a novel engine providing custom publishing capabilities for hybrid learning education.

The adaptive and end-to-end nature of the hybrid design is deeply grounded in the existing literature while providing an end-to-end experience that can be modified or optimized depending on any number of inputs, constraints, etc. The hybrid learning approach creates an ecosystem (Figure 1) where

[1]By seamless we mean in terms of both ease of use, and the lack of proprietary physical or technological requirements.

EICS'15, June 23 - 26, 2015, Duisburg, Germany

ACM 978-1-4503-3646-8/15/06$15.00
DOI: http://dx.doi.org/10.1145/2774225.2774837

printed textbooks coexist with digital content, with each enhancing the other, and providing users with a low-barrier method for switching between digital and print. While some systems have tried to combine print and digital [14], METIS creates a unique perspective allowing content and student to transition from digital to print and back without proprietary hardware, or cumbersome special paper. The digital environment does not attempt to replicate physical interactions (e.g. XLibris[22]), which is challenging and often ineffective to do [14]. Rather, METIS is designed to be amenable to a wide variety of factors including, but not limited to: personal customization, teacher management, source integration and recombination, reading and note taking on multiple surfaces/devices, and smooth transitions between print to digital and back to print (as location and needs change).

Based on lessons learned from a formative study of 523 students at SJSU as well as prior art, we formulate 5 design requirements for a robust system. These design requirements, in combination with the 4 requirements for Active Reading support, guide and scope the development of METIS. Once built, feedback on METIS was garnered through a series of focus groups involving 32 educators and students (at both high school and college levels).

RELATED WORK

E-books, tablets, and laptops have initiated a revolution in the education market [33]. Education IT has a worldwide 3.9% compound annual growth rate, 4.7% in Asia/Pacific, and 4.4% in the US[2] with spending (on both infrastructure and new classroom IT) estimated to grow from 16,884.7 ($M) in 2012 to 20,649.3 ($M) in 2018[3].

Mangen, et al. [18] performed an in-depth study of the comprehension differences between PDF reading and print reading in a naturalistic setting. They saw that comprehension suffered in the PDF condition, highlighting the potential challenges of scrolling, information recall due to lost contextual-spacial cues, forced linear reading, and difficulty of looking up annotations when taking an exam.

Research systems have strived to make CS education more accessible to women [20], demystify elementary math through interactivity [28], better explain discrete math [12], provide active learning for math [30, 31] and even teach propositional logic [16]. In the humanities, HCI and technology has been used to explore democracy [32], teach English [13, 29], and even play a role in special education [11, 5].

One particular avenue of research in the HCI community focuses on **Active Reading** (AR) [23], which "can be characterized by the greater demand it places on the reader and her media and tools [...] frequently [involving] searching, highlighting, comparison, non-sequential navigation, and the like" [33]. In general, AR involves four key activities: annotation, content extraction, navigation and layout [25]. Existing literature has explored AR [1], the process active readers go through [33, 21, 27], and the challenges active readers face when using technology [26]. Through these studies, technology solutions have been created to support AR. We specifically highlight XLibris, Papiercraft, and LiquidText.

XLibris[22] is an e-book software system for legal research. Users can annotate digital documents using a tablet and stylus, create clippings from digital articles, and organize them with additional freeform notes in a digital notebook. This paper-like approach gave users valuable features, but the authors [22] and others [33] suggest that the skeuomorphic metaphor limits flexibility when compared to paper.

Papiercraft [14] creates a novel system that brings paper-like interactions to a tablet + stylus, using printouts on special paper (Anoto paper) and a special physical pen (Anoto pen) to capture physical annotations. While technically supporting a hybrid (paper+digital) style of learning, Papiercraft was limited due to lack of multi-document annotation support, limited margins, and proprietary hardware/software (Anoto ecosystem), making it cumbersome and expensive to use.

LiquidText [34] exchanges the paper-like interactions of XLibris and Papiercraft for direct manipulation of content through touch interactions. Users can pull out text from a document, combine disparate section content together, and transform original documents by collapsing sections. However, in spite of the system's novel interaction and robust feature set, the layout is limited by screen real-estate compared to a desk and does not support users' need to be offline or make physical notes.

This work builds on the successes of XLibris, Papiercraft and LiquidText to allow Learners to easily transition from digital to print and back to digital without specialized hardware.

SCOPE AND DEFINITIONS

METIS is focused exclusively on the creation of an active learning tool to provide access to reading material, engagement with said material and the integration of digital and print media. It is important to note that METIS is *not* a full course management tool, as it provides no exams/quizzes, homework submissions, course slides, etc.

METIS is an end-to-end system, including a complete suite of components that supports the flexible hybrid model. While each individual component alone may not be new, it fulfills a set of design requirements and is explicitly grounded in an upfront study, HCI literature, or AR literature, yielding a robust and novel system. Some prominent components include the content workflow, where content is created, ingested, indexed, and managed; the book generation workflow, where content is laid out, watermarked, and print-ready; the recommendation engine, where external content like images, videos or articles are offered; and the consumption/flex-note workflow, allowing readers to transition from print to digital and back again (hybrid learning) via whatever channel they wish. Unlike prior work, METIS requires no proprietary or mandated hardware (e.g., specialized paper or pens).

To ensure clarity and precision, we define the following terms. **Article:** Text with interleaved images (e.g., a chapter

[2] Worldwide and U.S. IT Education and Training Services 2013-2017 Forecast, IDC #240681, April 2013

[3] Worldwide and U.S. IT Education and Training Services 2013-2017 Forecast Update (IDC #244914, December 2013)

in a printed text book). **Text Book:** A collection of articles, combined and ordered with a title, description and cover - used exclusivelyfer to a printed book. **Folio:** A collection of articles, combined and ordered with a title, description and cover image. We choose the term Folio over "book" so as not to combat preconceived notions of what a book can and cannot do, when using this digital reading platform.

INITIAL SURVEY

To motivate and guide this work beyond the existing literature and update design principles for current (2014) education environments, we deployed a perception survey of digital reading environments within an educational setting. Our aim to was guide and scope a system to meet current student needs.

The survey was advertised through an email distribution list, and posters around the campus of San Jose State University (SJSU). Of the 523 students who responded[4], age[5] and education was fairly diverse[6]. Participants were told that the first 500 respondents would receive $20 coupon towards the purchase of merchandise sold by HP.

The survey largely focused on Likert scale responses to the past usage and value of physical and digital reading technologies. We highlight student responses in this section, and discuss the results in the following section, concurrently with the design requirements for METIS.

Past Usage

We asked students about their prior experience with both physical and digital textbooks. 27.57% of students had used only print textbooks, 2.14% only digital, 67.18% had used both, and 2.14% claimed to have never used either. Of those that had used both, 19.06% preferred digital texts, 20.30% preferred using both together, and 60.64% preferred print.

Based on their experiences, students were asked Likert questions about the importance of different attributes (Table 1). Students who reported only using physical (digital) texbooks were asked exclusively about printed (digital) textbooks. Students who used both were asked about both technologies.

Qualitative Responses

We asked the 67.18% of participants that had used both digital and printed texts to briefly highlight what printed (digital) textbook features are missing in digital (printed) textbooks. The majority of responses about useful printed textbook features centered around the ease of note-taking or highlighting content, "flipping" back and forth between pages or sections, and offline accessibility. In contrast, key features of digital textbooks included easy searching, portability/access, copy/-paste and interactive content (e.g., videos). While these responses did not conflict with results found in prior work, they were valuable in focusing on the limitations/features to build

[4]42.9% Undergrad, 57.1% Graduate; 77.86% Female, 20.6% Male, 1.56% Undisclosed

[5]45.63% between ages of 18 and 25, 29.52% were 26-35, 17.67% were 36-50, 4.85% were over 50, and 2.33 % undisclosed

[6]0.78% Arts, 16.50% Business, 0.77% CS, 18.84% Education, 1.94% Engineering, 0.39% English, 30.10% L.I.S., 5.83% Sciences, and 24.85% "Other."

	Question	Mean	Std. Div
Printed	Ability to Learn	4.48	0.83
	Convenience & Access	4.40	0.94
	Cost	4.32	1.03
	Ability to Purchase	4.16	1.08
	Interactivity with Pen/Pencil	4.02	1.31
	Fewer Distractions	3.96	1.25
Digital	Convenience & Access	4.48	0.92
	Cost	4.43	0.99
	Ability to Learn	4.24	1.06
	Ability to Purchase	4.00	1.18
	Ability to Take Notes	3.85	1.31
	Physical Attributes of Device	3.51	1.44
	Interactivity with Other Software	3.50	1.41
	More Fun	2.28	1.39

Table 1: Technology's Important Attributes (Most to Least)

Likert responses with 5 being "Very Important" to 1 representing "Not Important"

and challenges to overcome, and highlighting that current perspectives towards physical and print have not shifted even as IT has become ubiquitous in education.

DESIGN REQUIREMENTS

Based on the related literature and our initial survey, we formulated a series of design requirements (DR) for building an effective and compelling digital textbook experience.

DR1 **Promote High Quality Physical Interactions:** Paper provides many rich interactions [26] that cannot (presently) be equally replicated and replaced digitally [22] such as free-form pen annotation (e.g., with no response delay), flipping pages, and spacial arrangement of content on a table[7][3]. Our study and prior work both show that students simply prefer paper to digital if forced to choose. We therefore support physical interactions, rather than poorly emulating them digitally.

DR2 **Facilitate Multiple Physical Use-Cases:** Much of the literature focuses on on Active Reading use case support in the digital domain. However, transforming digital material into physical is more than just "hitting the print button." Any robust system should support multiple types of printing, optimized for different use cases (reading, studying, note-taking). This need for robust physical interactions with reading material is not emphasized in the literature, and was explicitly uncovered in our upfront study.

DR3 **Make Digital Interactions Uniquely Digital:** Digital text can provide a rich set of content beyond that of printed books[35]. This includes low-level features (search, media, copy/paste), as well as higher level features (omplex cross-document indexing, nondestructive editing, automatic organization, NLP analysis). Any digital solution should highlight and take advantage of these uniquely digital features, rather than skeuomorphically emulate physical analogs in a digital environment [22, 33, 3] .

DR4 **Smooth Hybrid Learning & Transitions:** As both digital and print systems have distinct advantages, students should be able to easily transition between *digital and physical artifacts and then back to digital*[17, 9]. In this full-circle approach (Figure 1), once physical artifacts are generated, students can reconnect to the rich digital media. Choosing one modality or another should not lock the Learner in. Unlike prior work (e.g. [14, 17, 9]), these transitions should minimize any mandated infrastructure during these transitions (specialized paper or pens). Study participants highlighted the value of easy and convenient access in both mediums, implying that the transition between the two should be equally convenient and accessible. Further, [26] highlights the need for material to be optimized to the digital medium, rather than performing a 1-1 mapping.

[7]Some solutions (e.g. LiquidText[34]) support 2D spacial layouts, however, this is limited by device resolution. If the content is large enough to read, users cannot see all notes. Likewise, if all notes are visible on a tablet or laptop, users cannot read all the content well. This is in contrast to notes spread out/organized on a large table.

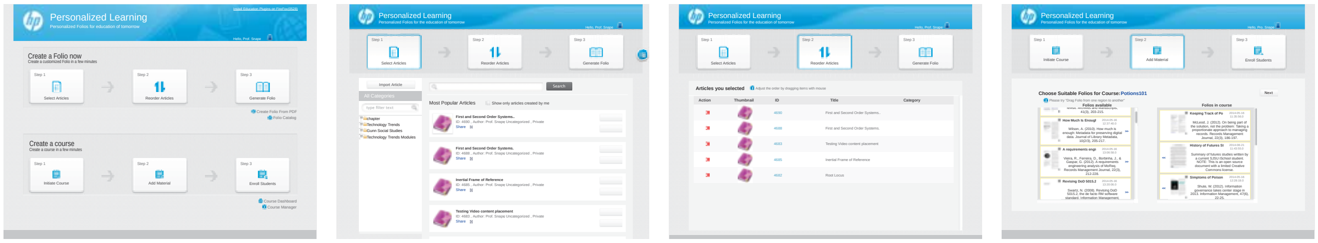

(a) Educator Homepage (b) Select and Upload Articles (c) Organize Articles for Folio (d) Add Articles to Folio

Figure 2: Educator Views

Images are high-resolution, can be zoomed in to a high degree on computer

DR5 Source Movement: Learners and Educators should be able to move across document sections and multiple documents, drawing connections between disparate elements and sources [26]. This flexibility is a hallmark of AR, and should not be limited to (or by) traditional sources. Specifically, this requirement focuses on the *type* and *source* of the material, **not** the method of movement between material.

DR6 Easy Folio Creation & Updates: From an Educators's perspective, a previously ignored area that is required as part of a rich education experience is the ability for Educators to easily generate Folios, updating their course material from a variety of sources with little effort.

DR7 Portability & Ease of Access: As highlighted in our study, portability and ease of access is an important aspect of any educational text, print or digital. A solution must function without Internet, allow computerless access (be printable without any special paper [14]), not require a large or cumbersome touch display (e.g. LiquidText[34]), and allow any artifacts (e.g., notes) or media (e.g., video) to be accessible in a physical *a*nd digital world.

In addition, the system must support AR Requirements:

AR1 Annotation: highlights and markings [23, 26]

AR2 Content Extraction: copying content for manipulation and organization elsewhere [23]

AR3 Navigation: "[M]ove through and between text via the most natural channel" (e.g., search, page turning, flipping between non-lineally connected text, etc.) [23].

AR4 Layout: arrange/compare material side-by-side, from any source[23]

This work strives to embody all AR support of the prior work, while providing seamless transitions between digital and physical material, physical interactions that can be fully captured and optimized for different use cases, digital environments that leverage computer indexing (multi-document cross-referencing, NLP based contextual searching, etc), and content creation/monitoring for Educators. Support of these features are novel in the context of the existing literature.

METIS UI/UX: FEATURES AND FUNCTIONS

Meaningful Education and Training Information System (**METIS**) is a novel software and service platform that provides a digital+print learning environment for faculty and students to manage and easily update course reading material, facilitate note taking, and monitor student engagement.

Each subsections' title highlights METIS compliance with a specific DR or AR, tying the engineering, interaction and design requirements together. To explicitly link our solution to the design requirements (and thereby the related work and study findings that generated the requirements) we cite the relevant DR and AR throughout.

The METIS system is based around a hybrid learning model (**DR4**). To facilitate this model and allow Learners to easily transition from digital to print and back to digital, METIS has two central features; Cheat Sheets & FlexNote, which allow for personal and customizable navigation (**AR3**) across multiple sources and documents (**DR5**). This hybrid learning solution leverages the best of paper[8] (**DR1**), and digital[9] (**DR3**). At the same time, the digital interface does not attempt to emulate physical interactions or the physical medium (**DR3**).

In many ways, METIS picks up where Papiercraft (2008) and XLibris (2001) left off, taking advantage of current technology trends which were not available when prior work was done. Where XLibris highlighted how powerful a system can be with AR support, Papiercraft hinted at the potential benefits of a paper-digital hybrid (though severely limited by their cumbersome proprietary hardware requirements). This work in turn looks to illustrate how AR and seamless hybrid-learning can be achieved through intelligent UX design and interactive-software engineering.

We illustrate the features and functionality of METIS through a scenario of Hermione, a student in Hogwarts and her teacher, Prof. Snape. Through their interactions with METIS, we will highlight the features, interface, interactions and use cases. This overly extensive scenario is meant to cover multiple approaches to interacting with METIS.

This illustrative scenario focuses on how existing techniques from AR research and other HCI systems can be built on and extended to create a truly seamless Hybrid Learning experience, centered around the use of printable (and digitally consumable) FlexNote and Cheat Sheet (without any proprietary hardware). While aspects of individual requirements may exist in prior work, METIS is unique in depth and completeness.

Educator: Creating Content (DR5, DR6)

As the semester begins, Prof. Snape brings up the METIS Educator interface (Figure 2) in his browser. Prof. Snape begins by uploading all the Articles for his Folio on Potions (Figure 2b). He pulls content from PDFs, URLs, Word Docs, and his

[8] Portability (**DR7**), easy pen-based annotation (**AR1**), and easy spacial arrangement and organization based on user needs (**AR4**)

[9] Rich media (**DR3**), cross-document content extraction/recombination (**AR2**) and easily flip (**AR3**) between multiple sources of content or search/navigate in multiple ways within/between documents (**DR5**).

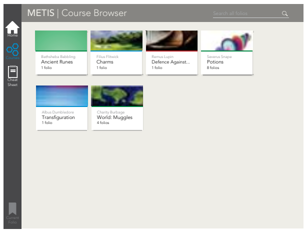

(a) Course Browser

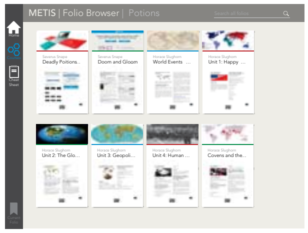

(b) Folio Browser

(c) Reading View

(d) Table of Contents

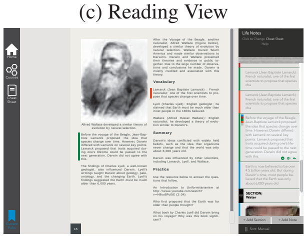

(e) Cheat Sheet

(f) Table of Images + Hover Menu

Figure 3: Learner Views

Images are high-resolution, can be zoomed in to a high degree on computer

own writing (**DR5**). He also adds other Educators' Folios or Articles to his Course, leveraging their expertise and reducing redundant work (**DR6**). Prof. Snape can create multiple Folios for each Course he teaches, that can be organized by unit, subject, source, or by calendar week.

Learner: Courses and their Folios

Once Prof. Snape sets up his Course and associated Folios, Hermione begins using METIS in her studies. Using her web browser, Hermione can browse Folios by Course (Figure 3a), author and recently viewed. Each Course (or author) is represented by a single card with Course, Educator, and image,[10] to act as a visual-memory cue [26].

An opened Course (or author) view displays all Folios (Figure 3b) with a combination of cover art (top quarter of the card), title, author and a screen shot[11] of the first page of the Folio.

Learner: Digital Reading (DR3, AR3)

Hermione next opens a Folio in the Reading View (Figure 3c). METIS provides Hermione with traditional page navigation options (**AR3**) from scrolling, a scroll bar (to the right of the Folio's content), by clicking on the left or right side of the Folio, or using arrow keys. To alleviate common scope problems in digital text [26], the scroll bar size provides a visual display of position relative to the length of the document.

[10] If an image is not chosen for the artwork, a color gradient is randomly assigned.

[11] Note only does this provide another visual landmark[26, 10], but is useful if a Educator uses the same cover art for all Folios in a given Course.

All content in the reader is generated using HTML, facilitating inline videos or interactive content inline with the text (**DR3**). This distraction-free reading environment allows Hermione to focus on the content (an important characteristic cited by study participants) without having to switch environments, causing a high cognitive load (a common digital environment problem[26]). The METIS interface utilizes a off-white page color to reduce eye strain, instead of a high-contrast black text on white field [19].

For a quick overview and navigation of the entire Folio, Hermione hovers on the left hand side of the screen, causing a translucent menu to appear (left side of Figure 3f). In this menu she can turn on a Table of Contents (left hand side of Figure 3d) or a table of images for more visual learners (left side of Figure 3f). These provide Hermione further methods for navigating her Folio (**AR3**), with the visual table of images providing visual landmarks [26] for navigation, contextual recall [10], and refinding content (**AR3**).

Learner: FlexNotes + Inline Cheat Sheet (DR5, AR1, AR2)

Like many other research systems focusing on AR (e.g., [33]), METIS provides Learners a mechanism for highlighting and noting important content, which we term the Cheat Sheet. Hermione can use the hover menu bar to enable/disable her inline Cheat Sheet (right side Figure 3e). When Hermione stumbles upon an important image or text to remember, she can highlight and drag the content to a Cheat Sheet, creating a **FlexNote** (**AR2** & **AR1**).

A FlexNote can contain any content (text or image). It is linked to a Folio fragment (paragraph, set of paragraphs, page), and has an associated color. When the FlexNote is created in the inline Cheat Sheet, a colored bar appears next to the FlexNote and to the left of the source's text/image in the body of the Folio. Thus, any FlexNote exists as a separate layer [26] from the text. By using this extraction+reference approach, Hermione can now expand, condense, or re-word the content of the FlexNote nondestructively (the original text in the Folio is not altered) making it easier to understand. Cheat Sheets also allow Hermione to organize her FlexNotes into sections by pressing the "Add Section" button (Figure 3e). FlexNotes can even be created without selecting content (**AR2** & **AR1**). When Hermione finds a whole page interesting, she can use the "Add Note" button in the Cheat Sheet (Figure 3e) and add free-form text to the Cheat Sheet, which will reference the current page.

FlexNotes are more than just a note-taking system. When Hermione clicks the color bar in the main text, the Cheat Sheet will open and scroll to the corresponding FlexNote (**AR3**). Likewise, the FlexNotes act as a personal navigation tool (**AR3**), and clicking on the return arrow on a FlexNote, will open the corresponding Folio to the origin of a FlexNote.

METIS takes the Cheat Sheet feature beyond prior work (e.g., [33]) by allowing the Cheat Sheet to span multiple Folios (e.g., one Cheat Sheet for all FlexNotes on a given subject mater), so Hermione can put new information next to older, creating a comprehensive and personalized view of subject matter across Folios, courses, and content sources (**DR5**).

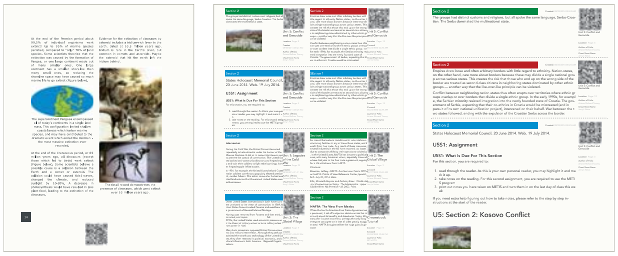

(a) Printed Page (b) Note Cards (c) Note Page

Figure 4: Printed Content

Images are high-resolution, can be zoomed in to a high degree on computer

Learner: Printing FlexNote (**DR1**, **DR2**, **DR4**, **DR5**, **DR7**, **AR4**)
The real power and novelty of the FlexNote and Cheat Sheet is exposed when integrated with the hybrid learning model (**DR4**). Hybrid learning gives Hermione the flexibility to choose how to study, which is very important according to both the related work and our study.

METIS makes it easy for Hermione to study for exams because Cheat Sheets & FlexNotes can be printed (**DR1**) in two formats (**DR2**): List View (a sequential list of notes - Figure 4c) and Flash Cards (a fixed box sized to an individual card - Figure 4b). Hermione never *needs* to worry about the source of the FlexNote (Folio A vs. Folio B) or that she printed a subset of her FlexNotes (**DR5**). She *can*, however, easily see a FlexNote's source, printed along the right side.

This easy-print rendering allows Hermione to cram for tests between classes (**DR7**), while walking to her dorm room (without having a laptop), or by spreading FlexNote out on a physical desk (**AR4**). In contrast, while other work has infinite 2D layouts for organizing digital notes (e.g.[33]), end users are still constrained by screen size, needing to balance readability of the note and ability to see all notes created.

The physical FlexNotes each have a QR code for quick indexing back to their digital counterparts. When the QR codes are scanned by her tablet, Hermione will jump directly back to the digital content and consume any interactive digital elements. The QR codes are unique to Hermione. The printed FlexNote cannot be easily be stolen because of the unique QR code would prevent any Learner who is not Hermione from indexing back to the source material. The unique QR codes also prevent cheating when Prof. Snape grades the FlexNotes (in many AR classrooms, Educators grade Learners' notes).

Learner: Recommendations (**DR3**, **DR4**, **DR5**, **AR2**)
When Hermione encounters a concept she doesn't understand, or wants to know more about, METIS leverages the technology from [15] to create an in-context search system (**DR3**). Hermione activates the hover menu and clicks the Search Icon (Figure 3f). This allows her to take any selected text and get a set of search results from resources (configured by the Educator, Wikipedia, YouTube, or any other standardized website). The approach in [15] allows METIS, and thereby Hermione, to search on a high level concept dynamically extracted from a highlighted paragraph or full pages, providing more information without a Educator present.

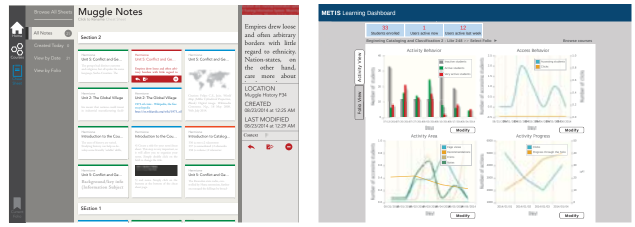

(a) Cheat Sheet View (b) Educator Dashboard

Figure 5: Additional Views

Images are high-resolution, can be zoomed in to a high degree on computer

This solution directly leverages the power of digital media (**DR3**) and addresses a major challenge: "[external material] exists in separate virtual spaces. Thus, drawing a new document into focus, or setting them aside, requires a shift of attention to a different representational space [... taking] more time and effort" [26]. With in-context recommendations and results, Hermione never has to leave her environment, switch windows, or be limited by artificially created boundaries between documents, and external resources (**DR5**).

If Hermione encounters noteworthy search results, she can drag them to her Cheat Sheet, making a new FlexNote for reference to later (**DR5**). Thus even digital content outside of her Folio can be captured, reproduced and organized both digitally and physically (**AR2** & **DR4**). When printed, these external digital references become physical, layering physical interactions on artifacts that were previously only digital. This further extends our Cheat Sheet and FlexNote infrastructure well beyond prior art.

Thus, even without Prof. Snape present, METIS' recommendation and FlexNote system allows Hermione to ask questions and get answers through METIS, pulling content from sources within and beyond the required reading (**DR5**).

Learner: Printing Folios (**DR1**, **DR2**,**DR4**, **AR1**)
Some institutions may choose to provide students with a fully printed copy of the Folio, while others may not. Regardless, METIS allows any Learner to print the digital Folio at any time (**DR1**). Even if Hermione forgets her physical text at home, she can print one page or one chapter from her Folio at any printer (**DR2**). These printed pages (unlike their digital counterparts in Figure 3c-3f) print with a QR code. Like a printed FlexNote, these QR codes, when scanned, will allow Hermione to jump directly back to digital content and consume any interactive or playable digital elements[12].

These QR codes provide further value. While reading the printed Folio, Hermione often makes notes in the margin of the pages, supporting her preexisting method of annotating and note taking (**DR1** & **AR1**). She can then take a picture with her phone and email the photo to a METIS service. No special application is need. This METIS service uses the QR code to identify the page and the user, and creates a new

[12]The authors are replacing the QR codes with HP LivePaper to create a hidden watermark (small very light fluctuations in the white background of the printed page) that uniquely identifies the page and the parts of the page. Using a mobile phone or tablet, an augmented reality experience is created with digital content appearing on the page itself, and jumping back to digital pages.

FlexNote from the physical notes. Next time Hermione opens METIS on her laptop, she will see her new FlexNotes, completing the digital to print to digital lifecycle.

Learner: Cheat Sheet & FlexNote Management (AR4)

Hermione has generated a large number of FlexNotes, from various sources and mechanisms. While the inline Cheat Sheet view is convenient for taking and extracting FlexNotes, the small screen real-estate makes in-depth organization challenging in the reading view, especially as the number of FlexNotes grow. METIS solves this problem similar to other research systems by treating Cheat Sheets and FlexNotes as first order data. Hermione has a dedicated Cheat Sheet view (Figure 5a) where she can see each FlexNote as a card, whose size scales with the amount of content contained within. Each FlexNote and sections are automatically laid out in the Cheat Sheet. From this view, Hermione can reorganize the FlexNotes and create new sections (**AR4**). Unlike prior work [34], METIS explicitly does not have a 2D layout management. In our upfront study, students valued the large spacial arrangement of notes on a table. While software can emulate an infinite 2D layout canvas for digital notes, small tablet and laptop screens hamper readability when "zoomed out," while losing the big picture when "zoomed in."

By clicking a FlexNote card, Hermione can edit the card's content with the edit window on the right side of Figure 5a. Editing in this view (or in the inline Cheat Sheet view) is full HTML, allowing for any markup and rich content (e.g., images or video). The edit window also allows Hermione to see where the FlexNote came from and when it was created, and even provides a contextual preview of the source above and below the original FlexNote.

Learner: Extensions and Offline Mode (DR7)

Sometimes during trips to study in the quiet cafe off campus, Hermione does not have Internet access. METIS provides Hermione with lightweight Chrome & Firefox extensions (**DR7**). When logged into METIS via her browser, she can use the extension to download the content, current progress, and all meta-data for any given Folio. Thus, even when offline, Hermione can have full access to all her study material. By using the browser sync storage, all of Hermione's enabled devices can maintain the same content (**DR7**).

Educator: Monitoring Engagement

Now that his Course is in full swing, Prof. Snape can monitor Learner engagement with the reading material. With any digital system, there is an important balance between information the technology *can* log, the information that *should* be logged, and the information that *needs* to be presented to a third party. METIS tackles this important privacy concern with its Educator Dashboard (Figure 5b). In the Dashboard, Prof. Snape can see the aggregate activity for his Learners across an entire Course and by Folio. He can observe engagement with the Folios, how much content is being recommended outside of the Folio, printing activity, and FlexNotes created. He has a strong sense of how much material Learners are engaged with, without making the Learners feeling like "big brother" is watching.

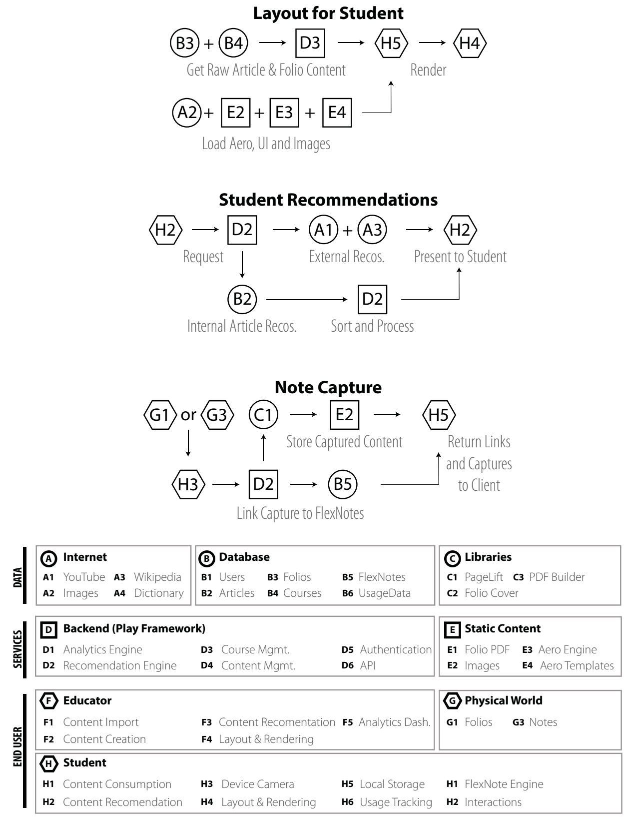

Figure 6: METIS Platform Design and 3 Example Data Flows

Image is High Resolution - Please Zoom Digitally in your PDF Reader

METIS ARCHITECTURE

The METIS platform is divided between server-side components and client-side components. Figure 6 shows the major METIS modules, as well as the data flow for three main features: content layout, note capturing and recommendations.

For content layout and rendering METIS uses a custom built engine called AERO, providing custom publishing capabilities. AERO takes as input content to be rendered and a set of templates. The content layout and rendering capability and data flow is depicted in Figure 6.

A METIS Folio is simply a collection of Articles in a predefined order. The content is composed and forwarded to the Learner device. The Learner device also fetches the AERO Engine code, written in JavaScript, and the AERO templates. Images or other multimedia elements used in the content are fetched from the METIS Static Content repository or the Internet. Using the templates, AERO generates a layout of the content by optimizing an aesthetic scoring function (e.g., by minimizing the amount of left-over white space left in the template [4] or minimizing the distance between text containing a reference to a figure and the figure itself [2].) The content and the code sent to the Learner's device goes through the Learner's Local Storage, which could potentially cache it for offline use or faster access in the future.

AERO runs entirely in the client browser. Using a greedy algorithm for each page, a template is picked from the suite and is filled with content chunks. The score is measured at each step and the best score, template and content combination is remembered, making long or complex Folios a trival extension of the base rendering engine. To optimize the rendering efficiency we save the output of the Folio, instead of using AERO to re-render the content every time.

For generating the PDF version of the Folio, we extend the PhantomJS library. PhantomJS acts as a headless browser and runs the AERO engine. Once the content is rendered (stripped of excess content and/or reformatted), PhantomJS saves the HTML output in PDF format.

Markup & Accessibility

By virtue of METIS Articles being based on HTML, content can contain text, images, video and other elements supported by HTML. As screen-reader accessibility is an important aspect of education, METIS denotes headings (`H1` to `H6` tags), figures (using the HTML5 `FIGURE` tag), and other markup. To ensure METIS accessibility, we conform to WAVE[13].

Facilitating Interaction with Digital Reader

For a better reading experience, AERO uses the zoom.js[14] library to adapt the page content to the current browser window width and height. This allows AERO to elegantly adapt to different screen resolutions as Learners switch devices and platforms (**DR7**), rather than "naively" scaling up content [26].

FlexNote Architecture

Each FlexNote is associated with a user, a Cheat Sheet (and section, if applicable), a Folio and page, a start and end element ID (e.g., paragraph or figure), and a start and end offset in the start and end elements respectively. A FlexNote also begins with the selected text in rich HTML, an assigned color bar color, a creation date, and a modification date.

The Learner can also capture notes made on the physical book or FlexNote using a cell-phone camera or a regular webcam. The data flow for this feature is depicted in Figure 6. Once the Learner captures an image, she sends the image to a predefined email address. The Content Management module receives the email and extracts the image. Next, the image is passed to the PageLift module which applies a suite of color and geometrical adjustments to the image in order to improve clarity. The technology behind PageLift has been described in [6, 8, 7] and solves many of the nontrivial challenges around image alignment, rectification, and enhancement. The enhanced image is stored as part of the Static Content. The PageLift module also performs QR code recognition and identifies the page or FlexNote. Finally the Content Management module combines the email address and ID in the QR code, and links the enhanced image to the Learner's Cheat Sheet, last used to take notes on the same book.

Recommendation Engine

In METIS, the Learner has immediate access to the content that the Educator has provided and the option to explore beyond what is provided, facilitated by the Recommendation Engine. Figure 6 shows the data flow for a recommendation request. The flow is initiated by the student selecting a content fragment and making a request to the Recommendation Engine. The Recommendation Engine fetches the relevant articles from the database and extracts the main topics from the content. It next performs a search on YouTube and Wikipeidia using their external APIs (additional content repositories could be easily added if they can be accessed through APIs). The recommendation results are then sent to the Learner's device. Finally the Learner's device fetches the recommended resources from their respective repositories.

Offline Support

We address the need for offline support by providing a browser add-on that is fully integrated with the online platform. The add-on is developed for Google Chrome and Mozilla Firefox browsers, leveraging the local storage in the browser extensions to download "offline" content. By using local storage in browsers that have cross-browser sync we are able to keep multiple browsers for a given user in sync with offline content and preferences.

FOCUS GROUP EVALUATIONS

In an attempt to validate our design decisions, identify areas of opportunity, and determine potential hurdles, we conducted focus group studies with Learners and Educators from both high schools and colleges. Four traditional focus groups were convened: College Professors/Instructors, College Students, High School Teachers, and High School Students. All groups contained individuals from multiple institutions, teaching/majoring in a variety of subjects (with the exception of high school students). Students varied in age[15], class year, and GPA. Each 8-person focus group participated in a 2 hour ssession (run by a contracted professional facilitator), consisting of a hands-on demonstration of METIS (with a discussion of features), followed by an in-depth round-table discussion to identify both sucessful features and areas for refinement.

There was nearly uniform consensus that METIS deliveres as a full end-to end hybrid learning model (**DR4**) for educators, and that students actually would want to use both digital (**DR3**) *and* print (**DR1**) features. As detailed below, each of the DR instantiations in METIS were praised. In a few situations, Learners and Educators wanted the technology to push even further, well beyond what our initial study and prior work suggested. The remainder of this section presents the focus group findings and their relationship with our original design requirements, thus connecting the literature/upfront-survey, design requirements, implementation, and evaluation.

Summary of Findings

When discussing physical/digital options, the most surprising was the disconnect between educators and students with regard to their needs and wants. Educators want students to read content online, post supplemental articles, and interact online (forums, chat, etc.). They also want powerful tracking of student activity, comprehension, etc. Students want to study material via print (even if distributed digitally) and

[13] http://wave.webaim.org/

[14] https://github.com/hakimel/zoom.js

[15] College mean age 19.36 years, HS mean age 15.45 years

communicate directly with educators; wish to read as little as possible; have no interest in "supplemental" material or online interaction; and do not want to be tracked at all. The design decisions of METIS appear to sit comfortably between both groups despite the impedance mismatch (control for educators and independence for students). Educators are able to link, collect and leverage any content they wanted (**DR6**) while students could consume the material through their desired medium (**AR3**, **DR4**) often via print (**DR1**) or in exported synthesized formats via printed Cheat Sheet (**AR2**). Our monitoring system gives educators broad course statistics about different types of engagement, while aggregating and anonymizing data to protect student privacy.

After being introduced to METIS, educator praise centered around the ability to customize Course material beyond traditional publishers (**DR6** & **DR5**), and the ability to constantly update material as class needs change. One educator stated:

> *My god, all I'm thinking is how I wish I had this 20 years ago for my statistics book. What I was doing then was typing it out... taking it down to the xerox place... and then the kids were really good at finding typos*

However, once the excitement abated, educators expressed concern about copyright infringement, especially when pulling in published material: *"What about just copying and pasting articles off the Internet with no sourcing? That is what I'm trying to stop my students from doing!"* This unforeseen concern and reaction resulted in educators stating the need for explicit citations (visible to students) and more in-depth consideration of the issue of ownership of compiled material. Any solution like METIS leveraging public and purchased content must mitigate copyright and ownership issues.

Students saw value in key features such as the Cheat Sheet and FlexNotes (**AR2**), and the freedom from carrying around a heavy book (**DR7**). Most positive comments focused on the ability to create flashcards and seamlessly consume them in hard copy (**DR4**, **DR2** & **AR4**). Students pressed for even more WYSIWYG controls (**AR1**) of FlexNote content: *"I like that I can make notes. Can I change the font colors? I color code my notes."* However, the customization needs for FlexNotes extended well beyond content and into more format options (the ability to customize printed card layout, visible information, source information, etc.)

Students felt that concept search was a good idea (**DR3**), but that resulting sources were too broad. Students wanted more focused concept search options, such as only "trusted" online resources/videos, or only reference sources, and were not satisfied with concepts from Wikipedia. Both students and educators praised the interactivity between paper and digital (**DR4**). The ability to use a smartphone to get at multi-media from a printed source or FlexNote was extremely powerful. Students expanded this to included language translation, and showing source or reference citations (to source information encountered in printed text).

DISCUSSION & LESSONS LEARNED

Overall, the focus groups confirmed the utility and value of our seven key design requirements as important needs/wants of Educators and Learners, though some resonated more with Learners, and others with Educators. Both saw high value in the hybrid learning model. We now wish to highlight three key additional lessons learned (in addition to our seven DRs) to help further inform the design of future systems.

LL.1 Print Media Is Not Dead — The value of print media is well established in education research, yet there is surprisingly little research on integrating print with digital content. Our work on METIS (and **DR1**, **DR2** and **DR4**), shows that the flexibility and ease of use was applauded and highlighted as the central great aspect of METIS for participants. Education solutions *must* smoothly transition content and interactions between physical and digital mediums.

LL.2 Allow Customizable Formatting — One unforeseen observation was the request by Learners to have even more control over the transition from digital to print, particularly in choosing the content, style and formatting of printed FlexNotes. This suggests an eighth design requirement for future systems, **Allow Customizable Formatting of Printed FlexNotes.**

LL.3 Content Matters — By easily cross-linking and examining content side-by-side, Learners as well as Educators really began to focus on having a rich set of content sources, which can be updated as needs change, and credited appropriately (see Future Work). Focusing on content implies the need to explicitly support it.

LIMITATIONS

While the final version of METIS was robustly built with a strong design process, there remain some limitations in our design. The offline mode presents a number of limitations of disabled functionality, with some being easy to resolve (creating FlexNotes, image capture, and book search), while others are very difficult to overcome (recommendations, voice reading, and dictionary search). The Educator side is not currently available in offline mode. The layout engine could use additional improvements as well. Currently a Folio is displayed the same way on any device, from mobile phone to workstation. We can consider device size when displaying content, while preserving layout consistency across devices and the physical book. The image capture engine, while quite powerful, has many possible improvements, such as more advanced detection of marked or highlighted regions, and interpreting types of annotations (arrows) rather than capturing the raw marks. Finally, our evaluation lacks a formal impact study.

FUTURE WORK

Based on our lessons learned, we envision two very explicit extensions of our METIS design. First (based on **LL.1**), we wish to further enhance the physical interactions and hybrid transitions with a more robust hand-written note capture system that would allow digitizing non-OCR markings such as arrows, underlines, or circles in conjunction with physical proximity on a printed page, to relate hand annotations to the source material. Second (based on **LL.3**), there was an open question posed by Educator in the focus groups as to plagiarism and rights management. This applies to both Educators adding content, and Learners bringing external content in (and crediting the source when they take FlexNotes.) Though outside our scope, we believe that these are important questions as digital platforms like METIS take shape.

We believe METIS struck a successful balance in the relationship between monitoring performance and student privacy.

However, this privacy concern also extends to collaboration, with students viewing a reading and note-taking experience as being a solo activity requiring privacy, while Educators see digital tools as platforms for interactive discussions. A more in-depth investigation should be conducted to be sensitive to the needs of all parties.

It is worth noting that there are still technological features we did not incorporate into the METIS system. For example, work like [34] leveraged the multi-touch platform with many novel interaction techniques that may fit well within METIS' digital experience. In this regard, we believe that our design requirements can be integrated within other published research systems to enhance their users' experiences.

CONCLUSION

METIS presents a complete digital to paper to digital hybrid learning framework for educational reading material. Unlike the existing literature, METIS allows Educators to build content and Learners manipulate content across multiple sources in the physical and digital world. The outcome of the entire Learners' interactions can seamlessly transition between media without any prerequisite hardware, paper or pens. "Printing" with METIS can support 3 distinct use-cases (reading, flash-cards, notes), while digital interactions are extremely robust, facilitated by complex NLP analysis of text and context. The engineering of this system to support the design requirements are unique, and while some aspects of individual requirements may exist in prior research, the depth and completeness to which METIS satisfies them is novel.

This rich set of interactions, centered around a Hybrid Learning Model, are explicitly facilitated with our METIS' novel flexible Cheat Sheet and FlexNote system for students to manage notes and navigate between multiple documents, on top of a custom digital document rendering engine called Aero. These features, and a handful of other minor contributions, are directly informed by a formative study of 523 students that generated a set of 7 design requirements. METIS followed onto these requirements as well as the 4 key interactions required for active reading. METIS was then validated through a series of 4 focus groups totaling 32 participants (Educators and Learners) from high school and college with participants explicitly validating the design decisions and METIS' unique functionality.

REFERENCES

1. Adler, A., Gujar, A., Harrison, B. L., O'Hara, K., and Sellen, A. A diary study of work-related reading: design implications for digital reading devices. In *CHI*, ACM Press/Addison-Wesley Publishing Co. (1998), 241–248.
2. Brüggemann-Klein, A., Klein, R., and Wohlfeil, S. On the pagination of complex documents. *Lecture Notes in Computer Science 2598* (2003), 49–68.
3. Chen, N., Guimbretiere, F., and Sellen, A. Designing a multi-slate reading environment to support active reading activities. *TOCHI 19*, 3 (2012), 18.
4. Ciancarini, P., Di Iorio, A., Furini, L., and Vitali, F. High-quality pagination for publishing. *Software—Practice & Experience 42*, 6 (June 2012).
5. Cramer, M., Hirano, S. H., Tentori, M., Yeganyan, M. T., and Hayes, G. R. Classroom-based assistive technology: collective use of interactive visual schedules by students with autism. In *CHI* (2011), 1–10.
6. Fan, J. Enhancement of camera-captured document images with watershed segmentation. *CBDAR07* (2007), 87–93.
7. Fan, J. Robust color image enhancement of digitized books. In *ICDAR'09*, IEEE (2009), 561–565.
8. Fan, J., Lin, Q., and Liu, J. An image processing pipeline for mobile document capture. In *ICIS2010* (2010).
9. Guimbretière, F. Paper augmented digital documents. In *UIST*, ACM (2003), 51–60.
10. Hailpern, J., and Huberman, B. A. Echo: the editor's wisdom with the elegance of a magazine. In *EICS*, ACM (2013), 313–322.
11. Hayes, G. R., Gardere, L. M., Abowd, G. D., and Truong, K. N. Carelog: a selective archiving tool for behavior management in schools. In *CHI*, ACM (2008), 685–694.
12. Henderson, P. B., and Marion, W. Interactive activities for learning discrete mathematics concepts: tutorial presentation. *Journal of Computing Sciences in Colleges 22*, 1 (2006), 65–66.
13. Kam, M., Ramachandran, D., Devanathan, V., Tewari, A., and Canny, J. Localized iterative design for language learning in underdeveloped regions: the pace framework. In *CHI*, ACM (2007), 1097–1106.
14. Liao, C., Guimbretière, F., Hinckley, K., and Hollan, J. Papiercraft: A gesture-based command system for interactive paper. *TOCHI 14*, 4 (2008), 18.
15. Liu, L., and Koutrika, G. Learningassistant: A novel learning resource recommendation system. In *Submitted to ICDE*, ACM (2015).
16. Lukins, S., Levicki, A., and Burg, J. A tutorial program for propositional logic with human/computer interactive learning. In *ACM SIGCSE Bulletin*, vol. 34, ACM (2002), 381–385.
17. Mackay, W. E., Pothier, G., Letondal, C., Bøegh, K., and Sørensen, H. E. The missing link: augmenting biology laboratory notebooks. In *UIST*, ACM (2002), 41–50.
18. Mangen, A., Walgermo, B. R., and Brønnick, K. Reading linear texts on paper versus computer screen: Effects on reading comprehension. *International Journal of Educational Research 58* (2013), 61–68.
19. Mansfield, J. S., Legge, G. E., and Bane, M. C. Psychophysics of reading. xv: Font effects in normal and low vision. *Investigative Ophthalmology & Visual Science 37*, 8 (1996), 1492–1501.
20. Marcu, G., Kaufman, S. J., Lee, J. K., Black, R. W., Dourish, P., Hayes, G. R., and Richardson, D. J. Design and evaluation of a computer science and engineering course for middle school girls. In *Proceedings of the 41st ACM technical symposium on Computer science education*, ACM (2010), 234–238.
21. Marshall, C. C., and Bly, S. Turning the page on navigation. In *JCDL'05*, IEEE (2005), 225–234.
22. Marshall, C. C., Price, M. N., Golovchinsky, G., and Schilit, B. N. Designing e-books for legal research. In *ACM/IEEE-CS joint conference on Digital libraries*, ACM (2001), 41–48.
23. Morris, M. R., Brush, A. B., and Meyers, B. R. Reading revisited: Evaluating the usability of digital display surfaces for active reading tasks. In *Horizontal Interactive Human-Computer Systems, TABLETOP'07*, IEEE (2007), 79–86.
24. Noyes, J. M., and Garland, K. J. Computer-vs. paper-based tasks: Are they equivalent? *Ergonomics 51*, 9 (2008), 1352–1375.
25. O'Hara, K. Towards a typology of reading goals rxrc affordances of paper project. *Rank Xerox Research Center, Cambridge, UK* (1996).
26. O'Hara, K., and Sellen, A. A comparison of reading paper and on-line documents. In *CHI*, ACM (1997), 335–342.
27. O'Hara, K., Smith, F., Newman, W., and Sellen, A. Student readers' use of library documents: implications for library technologies. In *CHI*, ACM Press/Addison-Wesley Publishing Co. (1998), 233–240.
28. Olive, J. Computer tools for interactive mathematical activity in the elementary school. *International Journal of Computers for Mathematical Learning 5*, 3 (2000), 241–262.
29. Rozan, M. Z. A., Mikami, Y., Bakar, A. Z. A., and Vikas, O. Multilingual ict education: Language observatory as a monitoring instrument. In *SEARCC*, Australian Computer Society, Inc. (2005), 53–61.
30. Sanchez Torrubia, M. G., Torres Blanc, C., and Krishnankutty, S. Mamadani's fuzzy inference emathteacher: a tutorial for active learning. *WSEAS transactions on computers 7*, 5 (2008), 363–374.
31. Sánchez-Torrubia, M. G., Torres-Blanc, C., and López-Martínez, M. A. Pathfinder: A visualization emathteacher for actively learning dijkstra's algorithm. *Electronic Notes in Theoretical Computer Science 224* (2009), 151–158.
32. Sarkar, M. S., and William, J. H. Digital democracy: creating an online democracy education simulation in a software engineering class. In *Proceedings of the 48th Annual Southeast Regional Conference*, ACM (2010), 55.
33. Tashman, C. S., and Edwards, W. K. Active reading and its discontents: the situations, problems and ideas of readers. In *CHI*, ACM (2011), 2927–2936.
34. Tashman, C. S., and Edwards, W. K. Liquidtext: a flexible, multitouch environment to support active reading. In *CHI*, ACM (2011), 3285–3294.
35. Thayer, A., Lee, C. P., Hwang, L. H., Sales, H., Sen, P., and Dalal, N. The imposition and superimposition of digital reading technology: the academic potential of e-readers. In *CHI*, ACM (2011), 2917–2926.