

Hybrid Documents Ease Text Corpus Analysis For Literary Scholars

Stephan Deininghaus Max Möllers Moritz Wittenhagen Jan Borchers
RWTH Aachen University
52056 Aachen, Germany
{deininghaus, moellers, wittenhagen, borchers}@cs.rwth-aachen.de

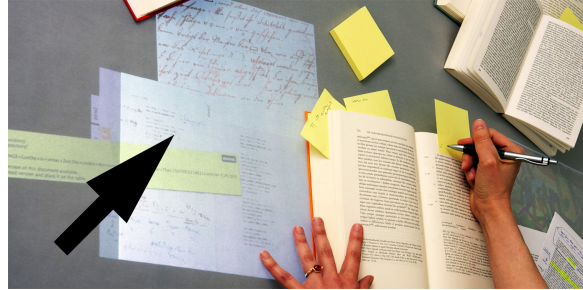


Figure 1: The designed system in use. Digital augmentation (marked by arrows) at the edge of a core physical document is integrated into expressive paper-centric processes on the desk.

ABSTRACT

We present a study that explores how literary scholars interact with physical and digital documents in their daily work. Motivated by findings from this study, we propose refactoring the working environment of our target audience to improve the integration of digital material into established paper-centric processes. This is largely facilitated through the use of *hybrid documents*, i.e., cross-modal compound documents that employ a printed book for rich, tangible interaction in tandem with a digital component for matching interactive augmentation on a digital workbench. The results from two user studies in which we evaluated increasingly detailed prototypes demonstrate that this design offers better support for central workflows in literary studies than currently prevalent approaches.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces.

General terms: Design, Human Factors

Keywords: Augmented workspaces, Digital workbench, Hybrid documents

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ITS'10, November 7-10, 2010, Saarbrücken, Germany.

Copyright 2010 ACM 978-1-60558-745-5/09/10...\$10.00.

INTRODUCTION

Digital documents have risen to widespread use as they offer considerable advantages over physical documents in some regards: transforming, searching, duplicating, and distributing them is comparatively simple. However, predictions that they would inevitably replace physical documents altogether turned out to be rushed, for printed documents continue to hold on to their edge in several important areas, such as richness of interaction and navigation, ease of annotation, mobility, and robustness [18, 22, 29]. Consequently, in knowledge work, it is common for many tasks to involve both.

Prior research has proposed numerous ways in which the affordances of paper and digital media could be combined to leverage their respective strengths to greater effect, for example by allowing documents to oscillate between physical and digital representations [10, 12] or capturing handwritten markings to provide enhanced capabilities to static paper [3, 6, 16, 17, 37]. In particular, one such approach has been to use the space provided by a tabletop for interactive augmentation of a physical document [15, 36].

Today, a common working environment for the knowledge worker features a desk for paper documents and one or more vertical displays for interaction with digital documents through a standard GUI. While such setups cater well to the aforementioned individual advantages of physical and digital documents, the inherent gap between the physical and digital realms hinders many forms of augmentation that call for close integration of modalities — basically, we are locked into two segregated work spaces.

This impedes tasks that involve both kinds of documents; in order to compare text in a book to text on the screen, the

reader may have to lift the book up next to the screen, or shift his gaze repeatedly between desk and display. Furthermore, the use of modality-specific document surfaces also precludes digital content from taking part in certain natural spatial layout strategies commonly found in working with multiple printed documents, e.g., placing related items close to each other in the workspace.

In this paper, we contribute a domain study that explores how literary scholars interact with physical and digital documents in their daily work. Armed with these findings, we then devise and justify a careful redesign of the working environment of our target audience to address identified problems. Thus, we show how tabletop-based document augmentation can be leveraged in real-world, multi-document knowledge work. We also evaluate our design using incrementally detailed paper and software prototypes.

RELATED WORK

For our problem domain, we could expect relevant prior research from philology (the study of language in historical sources) itself, from general research on reading and writing in multiple modalities, or from the evaluation of existing approaches to similar problems.

Philology Philological research has resulted both in extensive frameworks for the preparation and presentation of multifaceted text [25] and sketches of how information technology can offer novel angles on the text. Such visions exist for editing medieval manuscripts [8], expressing how a text changed over time [7], and creating a grid infrastructure for philological work [11].

Reading and Writing Research has explored the use and role of paper in knowledge work and general active reading processes in a wide range of contexts [1, 18, 23, 28]. Specifically, Sellen and Harper [29] argue that pen and paper will prevail in knowledge work because of their unique advantages.

Further fundamental research has widened the scope of these approaches by including the digital modality as well [22], recent work also addressing the distinctive affordances of novel displays such as tablets and interactive surfaces [20, 24, 31].

Implemented Systems Turning to applications of these principles, the characteristics of these processes have informed the design of systems such as *NiCEBook* [6] and the *Xlibris* active reading device [27]. The latter also saw adaptation to the needs of the professional domain of legal research [19]. Similarly, the *a-book* [17] and *ButterflyNet* [37] systems demonstrate how digitally captured physical annotation can be combined with other data to facilitate lightweight, expressive cross-modal input in other professional scenarios.

More generally, Wellner's *DigitalDesk* [34] has been the starting point of research on expansive, digitally augmented workspaces [14, 26] and interaction between tangibles and digital information [33]. As an extension of this approach, the *DocuDesk* system [10] by Everitt et al. demonstrated collocation of physical and digital objects on a shared surface to lessen the impact of modality on interaction.

With a more dedicated focus on interaction with single documents, both the *Interactive Textbook* [15] by Koike et al. and *WikiTUI* [36] demonstrated the merits of combining modalities in multi-component documents by augmenting the pages of a physical books with adjacent interactive content; *WikiTUI* also promoted the idea of using regular books as keys to retrieve associated digital information. Similarly, Steimle et al. [30] discuss the use of physical documents to control digital information on a tabletop, demonstrate how to integrate digital documents into original physical organizational processes on a single surface, and contribute findings on the occlusion of digital content by physical objects.

LITERARY CRITICISM

We now turn to literary criticism as the professional environment in which we take a closer look at these issues. Literary scholars are researchers that concern themselves with the analysis of text. The advantage of choosing this specific target audience is that it provides us with a professional application domain with a pronounced focus on interaction with text in which we can conduct studies and directly assess the validity of any designs resulting from these studies.

In particular, we are interested in the early stages of informal analysis, rather than more formal text production processes (i.e., we are interested in how the subjects would go about gathering information for writing an essay on a certain aspect of an author's work, rather than the stage in which they actually write that essay). When analyzing our application domain, we will pinpoint crucial differences to common knowledge work.

Initial Study

We conducted an initial study into the work processes of our target audience. It included semi-structured interviews with six scholars in both medieval and modern literary studies from the literature department of our university (three female, three male; from doctoral students to tenured academics). In addition to routinely performing textual analysis, all of them had professional experience with editorial projects, i.e., the selection and processing of material to be published for a scientific audience. We began these interviews by asking the subjects to verbally walk us through a typical part of their current daily work. We focused on descriptions of their workplaces, their interaction with printed and digital documents, and related note-taking habits. We then followed up on any points that we felt needed more detail, and discussed what aspects of working with paper and digital documents our subjects liked or disliked and how it influenced their work. With this understanding of the opinions and preferences of each participant, we concluded by letting them speculate freely about the nature of a "magical" device to assist them in their work regardless of current technical limitations. In addition, we asked one of the participants to perform an exemplary analysis task involving physical documents in a think-aloud session. We reviewed the footage from this session to retrace and refine the descriptions of work processes gathered from the interviews (cf. Fig. 2).

Working Environment

All of our subjects used a desk for paper-centric tasks in combination with a standard office computer for digital docu-

ments, online research, etc., much like it is the norm in other knowledge work as well.

Another characteristic property of the scenarios we evaluated is that the scholar will typically work alone, so that he can give his undivided attention to the task at hand. Thus, collaborative scenarios are rare during the actual analysis stage, and we will therefore focus on single-user tasks for the purposes of this paper.

Documents

Literary scholars require that the texts with which they work allow analysis and stand up to certain scientific standards, i.e., that they be “a reliable text that provides the basis for any historical or interpretative examination” complemented by an extensive collection of additional content (the *apparatus*) [25]. It is both the scientific rigor in preparing the text and the breadth and specialized nature of additional content — for example, a collection of textual variants, facsimiles of preserved original documents, commentary, historical background, or information on the text’s reception and transmission — that sets these *historical-critical editions* apart from regular reading editions and makes them a staple of the processes we examined in this study.

Since historical-critical editions seek to support analysis from a wide number of different angles, their contents are typically rather extensive to cover a large number of aspects; however, this also means that for many conceivable questions, much of that information might well be irrelevant. Another characteristic property is that much of an edition’s content is inherently descriptive, e.g., much of the commentary, variants, facsimiles, etc. in the apparatus describes (and, thereby, depends on) rather specific passages in other documents (often the main text). Thus, we can summarize by saying that editions exhibit a large degree of internal cohesion.

Traditionally, editions for philological work have been distributed exclusively as volumes of printed text. Along with the adoption of computers in knowledge work, however, some editions have since also been released — in part or whole — digitally, e.g., as a main printed component with supplemental content included on CD-ROMs, or as a full replacement for printed material such as web-based editions. Either digital content is typically accessed through standard GUI document viewing applications or web browsers on vertical displays. Today, all three kinds of editions are used in the field. As in general knowledge work, they are supported by other items such as conventional secondary literature, single articles, notebooks, scrap paper, or online content.

Structuring Access to Information

One central problem with printed editions is due to the inflexibility of paper. Oftentimes the amount of descriptive information that the editor wishes to provide on the main text will be too much to arrange with the main text itself. Therefore, such content is usually split off and provided separate from the main text. If the reader of the edition is interested in this information, he has to locate that information in the apparatus on his own and evaluate it in parallel to the main text. In deciding what information to include and how to structure the apparatus, the editor already has to act based on

assumptions on how the edition will be used; her choice may limit the edition’s usefulness for other conceivable scenarios. For example, certain editions enable the user to compare main variants of a text by presenting them side by side in adjacent columns (the technical term being *synoptic presentation*). Obviously, this makes comparing these variants much easier than having to reconstruct variants from an apparatus beforehand, unless the reader is interested in including a variant into the comparison that the editor has omitted.

Digital editions are more flexible at handling such problems by enabling the reader to customize the display of information for the current task. For example, some digital editions include facsimiles of handwritten documents that the user can choose to superimpose with a legible transcription. Furthermore, for editions that distribute the edited text digitally as well (i.e., fully digital editions), such information that supplements the edited text itself can even be inlined with it. This makes it unnecessary for the reader to connect and align edited text and descriptive content manually.

Additionally, digital content allows more advanced analysis. When we asked our subjects for advantages that digital documents have over printed documents for their work, full-text search came up most frequently overall. Thus, we were not surprised when this was also commonly named towards the end of the interviews, when we invited them to speculate about the “magical device” to aid them in their work. Besides search, other features from current digital editions were linguistic analysis and integration of content from external databases.

Functions of Reading

We observed that there appear to be at least two general modes involved in the study of text, which presumably serve different functions (cf. [1, 29]). Therefore, we will have to consider their influence individually.

First, we observed (and determined from the interviews) prolonged periods dedicated mainly to the absorbed sequential study of just one (or very few) documents. The focus of attention is generally stable in these phases. We conclude that the aim of these phases is to take in and process new information (“reading to learn” [1]).

On the other hand, we also witnessed phases in which attentional focus shifted frequently between several sources of information (i.e., documents, or possibly different passages of text within one document). Reading activity appeared to occur mostly in short, non-sequential spurs, attentional shifts often prompted by marks in the document (e.g., a cross-reference introduced by the editor, a short note scribbled in the margin, etc.). We assume that reading in these phases functions mainly as a trigger for cued recall of largely familiar content, rather than deep information processing [1]. We therefore conclude that these phases can best be characterized as serving the integration of information from multiple sources.

Single-Document Interaction

Research has shown that navigation in printed documents is aided by a number of automatisms. Both navigating to an approximate position anywhere in a book and turning pages




Figure 2: Typical configuration of documents in a literary scholar's workspace during a task that involves multiple documents in near-parallel (user seated at top edge of table).

when consuming content sequentially benefit from subtle haptic cues afforded by the document [18]. Furthermore, it has been argued that navigation in books can rely on shallow visual recognition processes, since the boundaries of the individual pages along with the specific shape of paragraphs, markings, etc. on them can provide a frame of reference with which we associate the actual content [22].

In fact, these findings apply to paginated digital content as well, so this is, at least in theory, a moot point; however, in our survey of the field, we found continuous, reflowed scrolling content to be the norm for digital editions. (See [9] for a discussion of alternative navigation techniques for digital documents.)

Managing Multiple Documents

From our study we concluded that it is a common task for literary scholars to act on information from multiple sources. For working with printed documents, this usually means having to arrange several items on a desk according to the specific task at hand. In our observations, we witnessed a high degree of reliance on spatial layout strategies when using multiple documents. For example, subjects laid out items in the workspace in such a way that documents with which they were currently concerned were placed at the near edge of the desk, while other items would reside in a kind of stand-by zone further away (Fig. 2). These configurations are effortlessly updated in response to attentional shifts and the requirements of the current task. For example, to compare text from multiple documents, these documents would typically first be placed close to each other to attain a makeshift synoptic layout. Additionally, spatial memory in such workspaces can provide strong recall cues in locating documents. Most of these observations are in line with prior research on physical document management in knowledge work [22, 23, 32].

This style of work allowed our subjects to take advantage of their spatial memory for establishing and updating document organization in their workspace. They could place documents on which they currently focused close to each other, near their position. For example, this is very useful when comparing documents, or taking notes in a notebook while reading

a book. Documents that they did not currently need they could move away, into a kind of stand-by zone. This scenario also affords direct manipulation, resulting in less cognitive overhead than mouse and keyboard-based interfaces on the vertical screen. In the interviews, our subjects also stated that they much preferred this style of work over digital document management because of these structuring processes and being able to “glance over” the documents with which they are working.

The setups consisting of vertical displays with standard GUI interaction that we considered in our study (i.e., as they are the norm for digital editions in today's literary criticism) do not offer comparable affordances for managing multiple documents. Consequently, no similarly rich spatial organization strategies were used in this modality. We attribute this partly to the additional complexity introduced by all manipulation having to pass through mouse and keyboard and partly on the limited screen real estate. These conclusions, too, parallel findings from prior studies on the topic [22].

Our interviewees also expressed their subjective preference for desk-centric document management over typical GUI interaction when working with multiple documents. Thus, our findings support prior research [22, 23] stating that such integrative tasks can benefit from spatial layout strategies.

Furthermore, the workspaces of desk and display are also permanently segregated, usually without even a connecting edge. This leads to another set of problems when a task involves documents from both modalities: digital documents cannot be integrated into those efficient spatial organizational strategies. For example, they cannot be part of a pile, and to compare paper and digital documents, the reader has to either hold the paper document up next to the screen, print the digital document (thereby manually transcending modalities), or shift his gaze between desk and screen repeatedly.

Note-Taking, Annotation, and Personalization

Our interviews showed that our subjects used a wide range of individual strategies for note-taking (e.g., writing on a piece of scrap paper, recording something in a text file, or in a private notebook) and annotation (e.g., underlining text in a book or scribbling in its margin). These activities were usually interleaved with the subtask at hand, regardless of how many documents that task involved. Interviewees also generally preferred note-taking and annotation in the context of paper-centric tasks over digital note-taking and annotation, citing richer expressivity for this — “the haptic joy of writing”, as one subject put it.

These findings agree with other results from general common knowledge work: paper affords richer, more effortless note-taking and annotation than digital documents, and generally affords closer integration of such active reading [28, 29]. Furthermore, handwritten marks on printed documents, by their virtue of being easily discernible from printed content, can support visual navigation in documents. In contrast, the perceived similarity between digital documents and user-supplied annotation is apparently one of the reasons why annotation of digital documents is often shunned [28].

One noteworthy point concerning our domain is that since users are often experts on certain topics, their personal library may contain documents that have been in use for a long time and have thus been personalized (both deliberately through markings, bookmarks, attached Post-It notes, etc., and implicitly, i.e., through stains or adaption of a book's binding to repeated strain) considerably, much to the benefit of their users (cf. [13, 16]).

Summary and Similar Domains

Our initial study showed that digital editions offer important functional advantages over physical editions, such as search and analysis facilities as well as highly flexible access to descriptive supplementary content, thus supporting a wide range of analytic angles on the text. On the other hand, we realize that physical documents remain common in these workflows because they afford efficient and natural interaction for a range of processes, from involved analytical study of single documents to integrating information from multiple sources. Thus, literary scholars have compelling reasons to accommodate both modalities in their work. In accordance with widely accepted observations [29], we predict that paper will continue to play an important role in literary work, and that our target audience will likely benefit from tighter integration of these modalities.

However, the observed processes are set apart from said common premises by a few characteristic differences. First, the amount of information is significantly higher and information is oftentimes more diverse, threatening to overwhelm the user. It is a challenge to provide means of structured access to this information. Second, the material exhibits pronounced interdependencies and internal cohesion since much of the content is both fragmentary and descriptive in nature and thus only meaningful in the context of the information it describes. This implies a hierarchy of dependencies within the material. Third, documents in personal libraries may see decade-long use, and researchers may have specialized in one author's work; both can lead to an exceptionally high degree of personalization and annotation accumulated in these documents.

In fact, we can use these characteristics to inform a search for professional domains in which similar challenges present themselves, and in which we might at least consider the applicability of our analysis and resulting design decisions. Trivially, general philology as the superset of literary studies might be considered a similar domain, building on the same editorial principles to convey information. Furthermore, material from legal research [19] shows at least some similar traits.

SYSTEM DESIGN

We propose a redesign of the common working environment to address the problems outlined in the previous section. Our design has two main foundations: an *interactive tabletop* to collocate physical and digital documents in a single workspace, and the adoption of a *hybrid document* model for tangible interaction with the documents themselves, which in turn enables the presentation and navigation of content via *panes*, *links* and *proxies*, and *excerption*.

Our design replaces desk and screen with a large interactive tabletop. We see three main advantages in giving our users a




Figure 3: Hardware schematics for the envisioned system. Three projectors project a 3240×1920 image at $135 \text{ cm} \times 80 \text{ cm}$ surface.

single unified workspace for both physical and digital items. First, readers can now organize digital documents using expansive spatial layout strategies — they can now arrange digital documents according to their current needs, see multiple digital documents in parallel, etc. just like with printed documents. Second, these digital documents now reside in the *same* space as physical documents, so there can be a single integrating organizational scheme. The user can place digital and physical documents side by side for comparison, arrange related documents close to each other in stand-by, place a piece of scrap paper directly next to a digital document when taking notes, etc. Third, this provides natural direct-touch manipulation. (While it is technically possible to use direct-touch interaction on vertical screens, this can become tiresome quickly.)

Our prototype tabletop (Fig. 3) has a surface area of $135 \text{ cm} \times 80 \text{ cm}$ at 3240×1920 pixels, resulting in a resolution of approximately 60 ppi. This is barely sufficient for reading small text, yet still below common desktop displays with about 100 ppi. Unfortunately, current projection-based tabletops generally offer low resolution and contrast compared to conventional screens. This affects reading performance and comfort negatively for digital text, especially since printed documents benefit from a well-lit environment. Thus, our design puts major focus on reading from analog media and only uses digital content for support and where printed media is not available.

Another problem is that of ergonomics, as reading text on a level tabletop requires hunching forward to get a good viewing angle. This puts significant strain on the reader's motor system. In contrast, physical documents can be tilted off the table, carried to a more comfortable location for reading, or repeatedly nudged in a largely subconscious process, thus cycling through different body postures [23]. However, the problem can be mitigated by tilting the desk at an angle towards the reader [20, 21].

Even with an unified workspace, the two classes of documents in it are still oblivious of their mutual existence. Furthermore, a new set of problems arises from co-locating modalities in this way. First, we are bound to run into occlusion issues,




Figure 4: Left: digital panes augmenting a physical book (ragged outline). The user can grab the individual panes by the provided tabs and show or hide them to customize the amount of displayed information for the core document. Right: an interactive cross-reference at the edge of a book and a proxy for a document that is currently unavailable.

although this problem may be overstated [30]. Second, when dragging digital content across the table, the user may have to maneuver around physical obstacles to get to his destination. Third, digital documents lack the materiality of printed documents that make room when shoved by another item (although there is research both on simulating this effect among digital items [2] as well as on automatic occlusion avoidance [5], we suspect it would be difficult to recognize the difference between accidental and intended overlap on a cluttered desk). And finally, we do not achieve full integration in the case of stacked documents (piles), as digital documents cannot be placed on top of physical items. Furthermore, although digital content on the tabletop can now be manipulated via touch — arguably an improvement over mouse and keyboard-type interaction — research indicates that this still does not afford quite the same richness of interaction as physical documents [31].

Therefore, in our design, we emphasize *hybrid documents*, by which we mean a cross-modality compound document that consists of a physical core document (e.g., in our domain, a printed book that contains the main edited text) and a supporting digital component that provides related information (e.g., commentary, textual variants, or references to secondary literature). This division is feasible because of the hierarchal and interdependent nature of much of the material in this domain, as outlined above. In order to display and move digital augmentation with its physical document, our design requires means to detect the identity, position, and currently opened page of documents on the table. We plan to achieve this through an overhead vision system, which can also provide scans of the documents. Numerous other methods for this exist [4, 10, 15].

The hybrid document model promotes physical documents to being tangible controls for digital content besides being first-class objects in their own right, much like Steimle et al. observed in a similar context [30]. The responsibilities of selecting (by turning pages) and positioning digital content are deferred to the physical document, reducing the number of degrees of freedom involved. Additionally, this decision mitigates the problem of having to slide digital content around physical obstacles when rearranging items in the workspace. It also helps relieve the reader of the burden of keeping reading positions in related physical and digital documents in sync manually.

Thus, our design achieves a favorable separation of concerns, in which we combine printed books — rich to interact with on and off the table, easy to read and annotate — with flexible digital augmentation. Note, however, that the design can accommodate purely digital documents trivially; it is that we simply *prefer* physical documents, when available.

Panes To interact with such augmenting digital information, the system displays tabbed, sliding *panes* alongside the edges of the core physical document (Fig. 4). The user can drag these out from under the document to display only information relevant to his current task. They are restrained to movement perpendicular to their respective edge of the document, with minimum and maximum extents. To illustrate the concept, suppose that the physical document is opened at a page with a poem. The system might then provide variants, commentary, and historical background for this poem grouped onto three panes alongside the document. If the reader is interested in seeing variants, she can expand the “variants” tab and review these side by side with the main poem, possibly even aligned at line granularity and with differences dynamically highlighted. To facilitate such alignment and make it more predictable how much space a pane will take up before expanding it, we recommend that they be no higher than the corresponding edge of the core document (for panes expanding left/right).

Overall, this design has the advantage of assisting in quickly selecting only relevant aspects from the the much larger amount of information typical of historical-critical editions. Since panes themselves move with their core document, the design introduces only a single additional degree of freedom per pane, thereby keeping interaction easy. Furthermore, panes expanding out to the left or right from the document are well-suited for the common case of presenting descriptive information aligned with the referenced main text, much like [15, 36]. Basically, the design adds a customizable set of additional columns of dynamic information to the printed document, thereby adapting well to both the requirements of the current task and available space on the desk — something that is often desirable but not feasible in printed historical-critical editions.

Links and Proxies One function that digital editions are often touted for is that of hyperlinking additional content. Our proposed design includes tap-activated links adjacent to printed text, or directly embedded in digital content. While

we cannot summon referenced physical documents to the table, we can do the next best thing: provide information on how to retrieve the document, or possibly display a digital version of the document if available. When such a link to external content is activated, and no digital version of the requested document exists, the system places a placeholder, or *proxy*, on the table (Fig. 4). We envision this proxy to be a free-floating bibliographic reference that is retained until it is explicitly dismissed by the user. It has to convey all necessary information to retrieve the represented document. When a referenced document is then placed on the table, it attaches to the outline of the document to indicate just the target page(s).

Proxies support a common subtask in active reading scenarios called *harvesting intentions* [16], i.e., going through a piece of text and collecting references to other potentially relevant documents for later examination. With the proxy concept, the reader can just activate promising links on the side while going through a document. When the reader finishes, there will be a collection of proxies for these referenced documents of which no digital representation is available. She could then operate on these items in a batch, e.g., print them all out, sweep them off the table onto a mobile device, forward them to the local library for retrieval, etc.

Excerpts Our design also supports the creation of full-page excerpts from physical documents using the overhead vision system. Similar to [10], this enables effortless transitions to the digital realm for further processing.

These basic building blocks are the framework for the organization, presentation, and navigation of hybrid content, of which we can compose advanced functionality. For example, a search function could include indexed physical documents, showing digitized previews of occurrences and using proxies to guide the reader; our software prototype uses panes to display related data from *TextGrid* [11], an initiative aimed at the support of philology.

EVALUATION

Before addressing such advanced functions, it is imperative that we first get a thorough understanding of how people interact with these building blocks. Thus, we conducted two qualitative user studies to evaluate the design, one with a low-fidelity paper prototype and another one with a software prototype for an interactive tabletop.

Paper Prototype

Study Setup Our primary goal in the first evaluation was to see if the design behaved as users expected, discuss its usefulness for our target audience, and identify usability problems. For this, a paper prototype (Fig. 5) was brought to life in *Wizard of Oz* manner, with the experimenter simulating system behavior. For the study, we created a book of poems in which we had replaced several central words with gibberish. Another book was a dictionary, containing the “translation” of the replaced words. We also prepared “digital” augmentation of the documents distinguished by differently-colored paper. This material contained textual variants and commentary (organized into sliding panes), and cross-references to other documents (both embedded in pane content and alongside the “physical” text). We also handed out a pen, a notebook, and a

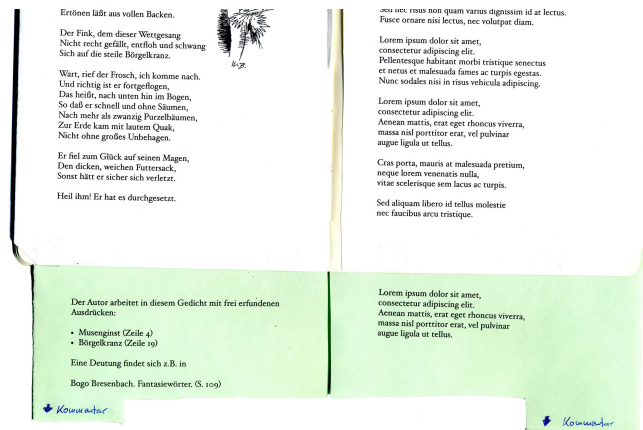


Figure 5: A part of the paper prototype, consisting of a “physical” document (top) with “augmentation” expanding downwards (bottom).

number of books that served as distractors. As a part of the instruction, we pointed out the distinction between “physical” and “digital” items, and asked them to verbalize their thoughts throughout the test. Snacks and drinks were placed on the same table to increase clutter. We recorded the interaction for later review.

In the test, we asked the subjects to locate a specific poem in the documents they had received and summarize its contents. To complete this task, they needed to explore the prototype, accessing both the augmenting panes and cross-referencing the dictionary to restore a meaningful version of the poem. This study design, which we had developed in cooperation with literary scholars, enabled us to mirror key processes from literary criticism (i.e., studying a text, retrieving variants and commentary from an apparatus, cross-referencing additional documents to integrate information) in a controlled setting. We also conducted exit interviews to address problems that occurred during the test, extract the users’ mental model of the system, get their opinion on individual aspects of the prototype, and discuss alternate designs. Sessions took about 45 minutes each.

Participants We recruited four students (all of them female, in their twenties) from the literature department of our university. All of them had recently completed a course on editorial work or were assisting in an editorial project, so that they had acquired sufficient basic training to qualify as members of our target audience. None of them had ever worked at an interactive tabletop before.

Results Accessing supplementary information through slide-out panes was well understood by all subjects. All of them assumed that the tabs at the edge of the book could provide additional information. On the first try, three subjects expanded a pane by tapping its protruding label, while one subject did this by dragging that label out from the book (both worked). No subject attempted to drag a pane past its maximum extension. Links displayed alongside the text and in digital material did not pose problems either, as all subjects assumed that these would point to additional information. All subjects activate them through tap. However, although all subjects

eventually managed to place the correct document on the table when cued to do so by a proxy object, all of them were surprised when it first appeared. Three subjects tried tapping it, and stated in the interviews that they had expected this to bring up either digital representation of the target document, or more information on how to retrieve it.

Physical occlusion of digital content occurred in all sessions, as the provided panes, when expanded, would frequently collide with other physical items on the table (in the case of expandable augmentation that was attached to the left and right edges of books) or with the near edge of the table (for panes attached to the bottom of books). We did not attach panes to the top of documents because preliminary tests revealed that these were partly obscured by the main document from typical viewing angles. We observed that subjects automatically rearranged occluding physical items on the table when collisions occurred to the left or right of an augmented document, but complained about collisions at the bottom, as this forced them to move the augmented book itself to an uncomfortable distance away from the edge of the table. We therefore decided not to pursue panes at the bottom of a book in future iterations. (Interestingly, [30] also remarks that subjects quickly relocated physical items to resolve occlusion, raising the question if the severity of this problem may be overstated.) Two subjects initially confused “physical” and “digital” items. This became clear when they attempted to pick up “digital” items to examine them more closely or relocate them to the other side of a document. Furthermore, in two sessions we observed a subject pick up a book from the table and recline for more comfortable reading.

Apart from the outlined issues, interaction was mostly flawless, and we deem the results from this first prototype evaluation promising. However, we feel that this study design only partially reflected the different affordances of paper and digital documents [31]. We did receive encouraging feedback from our subjects, who had themselves experienced problems concerning the integration of physical and digital content in literary criticism. All of them volunteered to participate in future studies and expressed interest in the progress of the system.

Software Prototype

Study Setup To validate and expand on the findings from the previous user study and follow up on the role of modality-specific affordances and challenges, we then built a software prototype and conducted a second user study on a rear-projected interactive tabletop (Fig. 1). We modeled the experiment after the paper prototype study in that we went for an exploratory setup with a think-aloud interaction part and a semi-structured exit interview. For the interaction tasks, we had prepared digital augmentation (alternative recipes, background information, cross-references to other recipes, etc.) for a physical cookbook and accompanying printed documents, in order to enable exploration, navigation and integration of information from multiple sources. We argue that this kind of material exhibits a structure and interdependency not unlike typical documents in literary work.

Due to technical concerns, we used an approximately 100 cm × 75 cm tabletop with a display area of 1024 × 768 pixels for

this study. Still, the surface was large enough to fit two physical documents with ample space for augmenting information and resolution sufficed for the text our material used. The height of the desk was between sitting and standing heights, facilitating both sitting use (with the user supported by a tall stool) and standing use for an extended reach envelope [32]. Standalone digital items, digital panes, links, proxies, and excerpts were fully functional via touch, and excerpting of entire pages was possible because we had prepared digital representations of relevant pages. Detection of document identity, position, and page number was not implemented in this version and we used another *Wizard of Oz* design to fake this data. We achieved very good reaction times to single or few consecutive page turns using keyboard shortcuts, whereas responding to documents being moved required more time since we had to match up the “frame” holding the digital augmentation by hand.

In the first task of the experiment, we studied the exploration of available information. To do this, we handed out a cookbook and asked our subjects to retrieve as much information as possible on a given dish. (The software augmented the printed core recipe with variants, historical background, and photos of the prepared dish as sliding panes and freely transformable digital photos that slid out from the pages.) In the second task, we studied navigation within and between documents. We pointed our subjects to an incomplete recipe in the main book and asked them to gather enough instructions for them to prepare that meal. This required them to navigate to content both within that same document and multiple passages in a secondary book. For the third task, we asked our subjects to pick a few dishes for a menu and clip information for a shopping list. This required them to excerpt and relate information from multiple sources.

Participants We ran the study with six students (four male, two female, all in their twenties), two from literary studies (one of which had already participated in the paper prototype study) and four from CS. While we did discuss the applicability of the design to literary studies with the participants from that field, the nature of the tasks in this study and its focus on general interaction with mixed-modality documents justify the participation of CS students.

Results Findings resembled those from the previous study in that accessing supplementary information in panes did not pose a conceptual problem for the participants; however, occlusion of expanding panes did occur in almost all sessions, typically caused either by pane content running off the table, or through other physical items. In line with our previous results, we observed that subjects quickly rearranged offending documents. Two subjects also expressed confusion when they expanded a second sliding pane on a page for which another pane was already expanded, which then caused the stack of expanded panes to run partially off the table. (Both subjects stated that they would have had expected the second pane to replace the first pane, not extend it.) Recognition and activation of cross-references to other documents worked satisfactorily again. However, several subjects reported that they had initially mistaken proxy objects for an error messages; one participant initially dismissed a proxy without reading

it. Another subject thumbed to the page number indicated on the proxy in the main document, failing to notice that it asked her to retrieve a different book. We were also surprised to see three subjects attempt to place the opened book face-down on the tabletop surface. They stated that they did this to make the contents of the book available to the system, much like using a scanner.

In the exit interviews, two subjects remarked that they would prefer to have some kind of explicit stand-by area for digital items, e.g., at the far edge of the table. We further note that the participants with a background in literary studies were particularly excited by the prospect of such an approach to modality integration for complex, multi-layered content. In fact, one subject expressed sincere disappointment that we had no plans to pursue commercial development.

Reception by Literary Scholars

The design was presented at an editorial philology conference, where it was met with great interest by members of the field. While there is a trend in our target audience towards the liberal and sometimes exclusive adoption of digital content, participants agreed that retaining the affordances of expressive paper-centric workflows by carefully integrating digital content into established desk-centric routines held great promise for conceptual tasks.

FUTURE WORK

We plan to conduct a long-term field study with a small set of professionals over the course of an entire philological project, i.e., from the early stages of analysis and conceptualization to the formulation of results. These projects typically run for multiple years or even decades, so this study should provide the possibility to assess how the design handles a full-scale real-world task, and if it influences the work of our users over time, similar to the study Wigdor et al. [35] conducted on the use of tablespots in a general office setting. What affordances our framework offers for the implementation of complex search or analytical functions is another interesting topic.

Furthermore, we are working towards a synergy with other current venues of research in computer-supported knowledge work. Note that in the analysis of domain-specific workflows, we concluded that rich support for annotation and note-taking was a key factor in efficient paper-centric workflows, which our design had to retain. Based on similar premises, related research has explored how physical handwritten notes can be captured by digital pens, paralleling our work in an effort to support natural processes with transparently deduced information [6, 17, 37].

Another promising area in this regard are approaches that aim to bridge the gap between physical and digital representations of one document [10, 12]. The workflows we analyzed in this paper typically comprise only a part of a larger creative process, the latter parts of which nowadays tend to center more around the computer-based production of original content. Therefore, providing a way for physical artifacts in the early conceptual stages to transition into a digital representation (i.e., digitization of content) better suited to these latter processes appears promising.

CONCLUSION

In this paper, we contribute a study of characteristic workflows in the analysis of text by literary scholars. The study explored how these workflows differ from other knowledge work in both process and the nature of involved documents, and highlighted the affordances of printed and digital content for these tasks. We then identified a set of fundamental challenges in this domain that are not optimally addressed by standard practices, which enforce a harsh separation of modalities. Our second contribution is demonstrating the design of a tabletop-based working environment that tackles this issue by closely integrating digital content into established desk- and paper-centric workflows using hybrid documents, in which physical core documents act as tangible controls for flexible digital augmentation.

Based on the implementation and qualitative evaluation of two increasingly detailed prototypes, we demonstrated that users were able to perform typical analytical tasks involving both physical and digital items with ease. Combined with encouraging feedback from professionals in the field, we conclude that this design framework has the potential to support core tasks in this domain better than current methods.

ACKNOWLEDGMENTS

This work was funded in part by the German B-IT Foundation and the Human Technology Centre of the RWTH Aachen University as part of the Excellence Initiative of the German Federal and State Governments.

REFERENCES

1. Annette Adler, Anuj Gujar, Beverly L. Harrison, Kenton O'Hara, and Abigail Sellen. A Diary Study of Work-Related Reading: Design Implications for Digital Reading Devices. In *CHI '98*, pages 241–248.
2. Anand Agarawala and Ravin Balakrishnan. Keepin' it Real: Pushing the Desktop Metaphor with Physics, Piles and the Pen. In *CHI '06*, pages 1283–1292.
3. Maribeth Back, Jonathan Cohen, Rich Gold, Steve Harrison, and Scott Minneman. Listen Reader: An Electronically Augmented Paper-Based Book. In *CHI '01*, pages 23–29.
4. Maribeth J. Back and Jonathan Cohen. Page Detection Using Embedded Tags. In *UIST '00*, pages 159–160.
5. Peter Brandl, Jakob Leitner, Thomas Seifried, Michael Haller, Bernard Doray, and Paul To. Occlusion-Aware Menu Design for Digital Tabletops. In *CHI '09*, pages 3223–3228.
6. Peter Brandl, Christoph Richter, and Michael Haller. NiCEBook: Supporting Natural Note Taking. In *CHI '10*, pages 599–608.
7. Dino Buzzetti and Malte Rehbein. Textual Fluidity and Digital Editions. In Milena Dobreva, editor, *Text Variety in the Witnesses of Medieval Texts*, pages 14–39, 1998.
8. Jonas Carlquist. Medieval Manuscripts, Hypertext and Reading. *Visions of Digital Editions. Lit Linguist Computing*, 19(1):105–118, 2004.

9. Andy Cockburn, Amy Karlson, and Benjamin B. Bederson. A Review of Overview+Detail, Zooming, and Focus+Context Interfaces. *ACM Comput. Surv.*, 41(1):1–31, 2008.
10. Katherine Everitt, Meredith Ringel Morris, A. J. Bernheim Brush, and Andrew D. Wilson. DocuDesk: An Interactive Surface for Creating and Rehydrating Many-to-Many Linkages among Paper and Digital Documents. In *TABLETOP '08*, pages 25–28.
11. Peter Gietz, Andreas Aschenbrenner, Stefan Budenbender, Fotis Jannidis, Marc W. Kuster, Christoph Ludwig, Wolfgang Pempe, Thorsten Vitt, Werner Wegstein, and Andrea Zielinski. TextGrid and eHumanities. In *IEEE E-SCIENCE '06*, page 133.
12. François Guimbretière. Paper Augmented Digital Documents. In *UIST '03*, pages 51–60.
13. William C. Hill, James D. Hollan, Dave Wroblewski, and Tim McCandless. Edit Wear and Read Wear. In *CHI '92*, pages 3–9.
14. H. Koike and M. Kobayashi. EnhancedDesk: Integrating Paper Documents and Digital Documents. In *IEEE APCHI '98*, page 57.
15. Hideki Koike, Yoichi Sato, Yoshinori Kobayashi, Hiroaki Tobita, and Motoki Kobayashi. Interactive Textbook and Interactive Venn Diagram: Natural and Intuitive Interfaces on Augmented Desk System. In *CHI '00*, pages 121–128.
16. Chunyuan Liao, François Guimbretière, Ken Hinckley, and Jim Hollan. Papiercraft: A Gesture-Based Command System for Interactive Paper. *ACM Trans. Comput.-Hum. Interact.*, 14(4):1–27, 2008.
17. Wendy E. Mackay, Guillaume Pothier, Catherine Letondal, Kaare Bøegh, and Hans Erik Sørensen. The Missing Link: Augmenting Biology Laboratory Notebooks. In *UIST '02*, pages 41–50.
18. Catherine C. Marshall and Sara Bly. Turning the Page on Navigation. In *JCDL '05*, pages 225–234.
19. Catherine C. Marshall, Morgan N. Price, Gene Golovchinsky, and Bill N. Schilit. Designing e-Books for Legal Research. In *JCDL '01*, pages 41–48.
20. M. R. Morris, A. J. B. Brush, and B. Meyers. Reading Revisited: Evaluating the Usability of Digital Display Surfaces for Active Reading Tasks. In *TABLETOP '07*, pages 79–86.
21. Christian Müller-Tomfelde, Anja Wessels, and Claudia Schremmer. Tilted Tabletops: In Between Horizontal and Vertical Workspaces. In *TABLETOP '08*, pages 49–56.
22. Kenton O'Hara and Abigail Sellen. A Comparison of Reading Paper and On-Line Documents. In *CHI '97*, pages 335–342.
23. Kenton O'Hara, Alex S. Taylor, William M. Newman, and Abigail Sellen. Understanding the Materiality of Writing from Multiple Sources. *Int. J. Hum.-Comput. Stud.*, 56(3):269–305, 2002.
24. Anne Marie Piper and James D. Hollan. Tabletop Displays for Small Group Study: Affordances of Paper and Digital Materials. In *CHI '09*, pages 1227–1236.
25. Bodo Plachta. *Editionswissenschaft: eine Einführung in Methode und Praxis der Edition neuerer Texte*. Reclam, Stuttgart, 1997.
26. Jun Rekimoto and Masanori Saitoh. Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments. In *CHI '99*, pages 378–385.
27. Bill N. Schilit, Gene Golovchinsky, and Morgan N. Price. Beyond Paper: Supporting Active Reading with Free Form Digital Ink Annotations. In *CHI '98*, pages 249–256.
28. Abigail Sellen and Richard Harper. Paper as an Analytic Resource for the Design of New Technologies. In *CHI '97*, pages 319–326.
29. Abigail J. Sellen and Richard H.R. Harper. *The Myth of the Paperless Office*. MIT Press, 2003.
30. Jürgen Steimle, Mohammadreza Khalilbeigi, and Max Mühlhäuser. Hybrid Groups of Printed and Digital Documents on Tabletops: A Study. In *CHI EA '10*, pages 3271–3276.
31. Lucia Terrenghi, David Kirk, Abigail Sellen, and Shahram Izadi. Affordances for Manipulation of Physical Versus Digital Media on Interactive Surfaces. In *CHI '07*, pages 1157–1166.
32. Aaron Toney and Bruce H. Thomas. Considering Reach in Tangible and Table Top Design. In *TABLETOP '06*, pages 57–58.
33. John Underkoffler and Hiroshi Ishii. Urp: A Luminous-Tangible Workbench for Urban Planning and Design. In *CHI '99*, pages 386–393.
34. Pierre Wellner. Interacting with Paper on the DigitalDesk. *Commun. ACM*, 36(7):87–96, 1993.
35. D. Wigdor, G. Perm, K. Ryall, A. Esenther, and Chia Shen. Living with a Tabletop: Analysis and Observations of Long Term Office Use of a Multi-Touch Table. In *TABLETOP '07*, pages 60–67.
36. Chih-Sung (Andy) Wu, Susan J. Robinson, and Ali Mazalek. Turning a Page on the Digital Annotation of Physical Books. In *TEI '08*, pages 109–116.
37. Ron Yeh, Chunyuan Liao, Scott Klemmer, François Guimbretière, Brian Lee, Boyko Kakaradov, Jeannie Stamberger, and Andreas Paepcke. ButterflyNet: A Mobile Capture and Access System for Field Biology Research. In *CHI '06*, pages 571–580.