HoloDoc: Enabling Mixed Reality Workspaces that Harness Physical and Digital Content

Conference Paper · April 2019

DOI: 10.1145/3290605.3300917

HoloDoc: Enabling Mixed Reality Workspaces that Harness Physical and Digital Content

Zhen Li
University of Toronto
Toronto, Ontario
Canada
zhen@dgp.toronto.edu

Michelle Annett
MishMashMakers
Newmarket, Ontario
Canada
michelle@mishmashmakers.com

Ken Hinckley
Microsoft Research
Redmond
Washington, United States
kenh@microsoft.com

Karan Singh
University of Toronto
Toronto, Ontario
Canada
karan@dgp.toronto.edu

Daniel Wigdor
University of Toronto
Toronto, Ontario
Canada
daniel@dgp.toronto.edu

ABSTRACT

Prior research identified that physical paper documents have many positive attributes, for example natural tangibility and inherent physical flexibility. When documents are presented on digital devices, however, they can provide unique functionality to users, such as the ability to search, view dynamic multimedia content, and make use of indexing. This work explores the fusion of physical and digital paper documents. It first presents the results of a study that probed how users perform document-intensive analytical tasks when both physical and digital versions of documents were available. The study findings then informed the design of HoloDoc, a mixed reality system that augments physical artifacts with rich interaction and dynamic virtual content. Finally, we present the interaction techniques that HoloDoc affords, and the results of a second study that assessed HoloDoc's utility when working with digital and physical copies of academic articles.

CCS CONCEPTS

• **Human-centered computing → Mixed / augmented reality**; *Interaction techniques.*

KEYWORDS

Mixed reality; Augmented reality; Digital pen input; Reading behavior

ACM Reference Format:
Zhen Li, Michelle Annett, Ken Hinckley, Karan Singh, and Daniel Wigdor. 2019. HoloDoc: Enabling Mixed Reality Workspaces that Harness Physical and Digital Content. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019), May 4–9, 2019, Glasgow, Scotland UK*. ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3290605.3300917

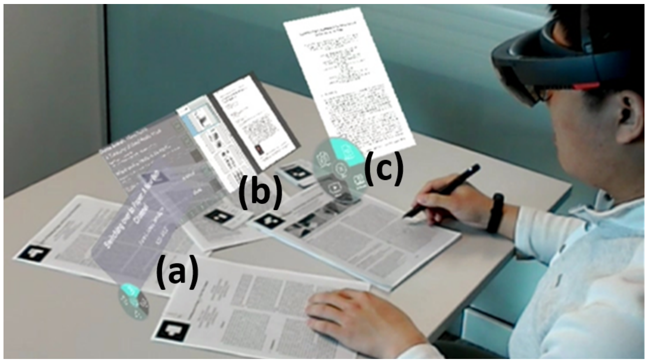

Figure 1: The overall view of HoloDoc where the user is able to (a) check meta-data information, (b) browse online search results, and (c) read related reference documents.

1 INTRODUCTION

With the advent of every new (mobile) computing technology and input interface design, we come one step closer to realizing the paperless office [52]. Although one would expect such technologies to decrease the volume of physical documents in workplaces, the volume of documents has increased in the decades following the advent of the personal computer [13, 52]. Significant research has investigated possible explanations for this, by examining the benefits of the physical attributes of paper in comparison to digital displays [3, 8, 42, 44, 57], such as the tangibility and navigability of paper [57, 59]. Others have demonstrated the superiority of physical documents for tasks involving large numbers of documents, such as for sensemaking [51]. Physical paper documents naturally support multi- and intra-document layouts,

---

CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

ACM ISBN 978-1-4503-5970-2/19/05…$15.00
https://doi.org/10.1145/3290605.3300917

in addition to providing simplistic readability and supporting physical gestural navigation. Notably, many of the studies on paper and digital displays predate the adoption of modern, high resolution displays and mobile devices by office workers, which could address many of these issues [8, 42].

Modern office workers have access to a plethora of digital devices, as well as cloud services and other technologies that can ease the movement of information between devices. In the present work, we first sought to update our understanding of the use of physical versus digital documents, to determine if modern tools have solved many of the challenges highlighted by previous findings. Further, to understand whether the results from the literature would hold for younger, technologically adept users who were equipped with modern technology and performed document-intensive tasks, a design study was conducted. This study utilized standard data analytic task, in which participants were asked to review over 200 documents [11]. The intent was to understand how users would use physical documents in light of easily accessible digital versions of said documents. From the observations, we found that despite advances in human-computer interfaces, participants still valued physical documents over digital versions for many tasks. This was due to physical documents' quick access, physical navigation, convenient markup abilities that persisted after the document was no longer the focal file, and their amenity to dynamic physical layouts.

Inspired by these results, we created *HoloDoc*, a mixed reality system that allows users to make use of physical documents, while preserving the advantages of digital content, such as the ability to search, hyperlink, and access rich media. HoloDoc tracks the location of physical artifacts, as well as a user's gestures and strokes, to augment the artifacts with multimedia content in Mixed Reality (Figure 1). Users are then able to directly interact with content in 3D space and not limited by the field of view of a camera or projector. A major contribution of this work is that HoloDoc implemented several features that fall along the "reality to virtuality continuum" [38]: augmenting physical artifacts with dynamic functions, attaching virtual elements to physical proxies, and managing virtual content that was transferred from the physical world. A second study was conducted to assess the usefulness of HoloDoc within traditional academic reading tasks. The results determined that participants who were not familiar with the concept of mixed reality could easily understand HoloDoc and were able to propose uses for the system within their daily research-based activities.

2 RELATED WORK

Of most relevance to the present work is literature relating to the tandem use of digital and physical documents, the augmentation of physical documents with digital content, digital pen-based systems for use with physical documents, and prior research on reading with physical and digital documents.

Mixing Digital and Physical Documents

Tablets and phones have been used to enhance reading activities by displaying digital documents or enabling additional functionality, such as the ability to search for additional content, provide a space for note taking [5, 17], or provide a "magic-lens" interface to visualize additional information about a physical document [16, 29, 63]. Interactive tabletops have also been used to augment the spaces surrounding physical documents [9, 34, 35]. DocuDesk, for example, augmented physical paper with shadow menus that were displayed around sheets of paper on a tabletop computer [9]. The user could then use the tabletop to explore the links between the paper and digital documents and further manipulate their relationships using a digital pen.

These approaches augment physical documents with rich content, however, visual feedback was limited by the dimension of the devices and augmented content was visible by passersby. With HoloDoc, not only are the digital augmentations user-specific, to not distract others who may be in the area, but they travel with the user's field of view, ensuring that they remain contextually and situationally relevant at all times. HoloDoc also enables users to leverage their spatial memory to retrieve tasks while managing windows in 3D space [26].

Augmenting Physical Documents

Many projects use stationary and mobile projection to augment physical documents. For example, the DigitalDesk used a projector and several cameras pointing towards the surface of a desk to track a user's finger movements and pen stroke, and displayed virtual content to augment the physical documents on the desk [60, 61]. Live Paper identified the cards, papers, and books on a desk using a camera, and projected virtual menus around these artifacts [48]. WikiTUI augmented physical books with rich resources from wikis [62]. The EnhancedDesk used an additional pan-and-tilt camera and printed matrix codes to enable better tracking of users' hand gestures, and provided various interactions to augment a textbook [25]. The Everywhere Displays Projector presented virtual content "in context" on a whiteboard, close to a phone, on top of a cabinet, or on a wall [45]. Hand-held augmented reality displays have also been explored. With MagicBook, a hand-held binocular display and AR tags printed on a book could be used to explore virtual worlds [2]. The Mixed Reality Book harnessed the affordances of a traditional book, while presenting virtual content in the book pages and the space around the book via a handheld display [10]. The

ARToolKit [22] has also spurred the development of paper-based augmented reality systems such as AR Lamp [23] and PaperWindows [18].

Unlike HoloDoc, these approaches restricted users to interacting with content in a specific area on the surface of a desk, with this active region limited by the position of the cameras and projectors that were used. HoloDoc presents user-specific, in-context virtual content that augments the physical environment in 3D space and enables users to utilize off-surface information displays via direct interaction.

Pen-based Interfaces

Pens are the most common peripherals used in conjunction with paper. They are not only used for writing, but also for pointing and marking on pages [31] and for visual tracking and annotating [19]. Studies by Riche et al. demonstrated that digital pens provide many affordances such as the ability to archive and retrieve handwritten notes [47].

Several systems have utilized digital pens with physical documents. Norrie and Signer presented the Paper++ architecture, which integrated printed and digital media for the user [31, 41]. By tracking digital pen movements on physical paper, associated multimedia files and hypermedia documents could be displayed on a tablet or monitor. PapierCraft introduced a pen gesture-based command system which leveraged the input capability of a digital pen for use with physical paper [28]. In addition to enabling pen gestures within a physical book, PAB used a paper GUI 'palette' to support mode switching [6]. PenLight [54] and MouseLight [55] used a mobile projector attached to a pen or mouse to display virtual content on top of paper so that users could interact with the displayed visualizations. NiCEBook combined the flexibility of taking notes on physical paper with the benefits of digital presentation [4]. IllumiPaper proposed a UI framework and prototype for interactive paper, focusing on visual feedback position, feedback time, and feedback types [24]. Steimle also proposed a framework for pen-and-paper user interfaces, and contributed several novel strategies for pen-based linking and tagging [56].

In HoloDoc, both digital pen strokes and hand gestures are used as input, i.e., the user invokes commands and corresponding virtual windows by tapping the pen on physical paper, and then performs further operations using hand gestures. In this way inking and command execution is separated, similar to MouseLight [55], while ensuring that the paper is free from additional pen strokes.

Reading Behavior

Much research has focused on understanding the benefits of reading on paper [8, 42, 44, 57]. Takano et al., for example, found that reading on paper was faster and had a higher error-detection rate than reading on a screen [57]. Work by Pearson et al. reported other benefits such as the creation of placeholders for document revisitation, annotations, note-taking, and visual indexing [44]. According to O'Hara and Sellen, annotating on paper while reading ensures a better understanding of content [42]. Bondarenko and Janssen identified that people print digital documents from various sources not only to read and annotate them, but also store them in *"one, easily accessible place"* [3].

Digital text, on the other hand, also has advantages: the ability to search [8], the ease with which content can be archived [13] or edited [59], and so on. Digital documents can also convey more dynamic information, for example, by using an animated figure rather than a sequence of static figures [12]. Although annotating on paper used to be a common practice, advancements in computing devices have improved so much that Morris et al. found users ranked tablets and paper similarly for reading and annotating tasks [39]. They also found that users liked having the freedom to rearrange tablets, but challenges such as insufficient margin space and an inability to easily move information between devices existed. To achieve the best of both worlds, HoloDoc uses a mixed reality approach, fusing physical paper reading experiences with the benefits of digital technology.

3 DESIGN STUDY

A design study was conducted to understand how modern users, who were equipped with their own workspaces and allowed to use any of their hardware and software, would make use of paper and digital documents when faced with a complex, document-intensive task. The study extends previous interview-based work [49] with direct observations of participants' own workspaces, and investigated participants' interactions with documents in both digital and physical forms, as opposed to only focusing on desk space organization [33] or computing with multiple devices [7]. By using participants' own environments, behaviors were not biased by being in an unfamiliar environment with unusual devices or fixtures.

Participants

Eight Computer Science student researchers were recruited to participate in the study (5 male; Mean = 24 years, Range = 18-31 years old). All participated at their own workspaces, using their own equipment, which included a desk, a chair, and at least one laptop or desktop computer (with a monitor). Three participants also had a tablet. None of the participants had experience with similar document analytical tasks. Participants were provided with $20 CAD as an honorarium for their time.

The choice to use students limits the potential generalizability of the results, however, this group is interesting to study because they are innovative and are regularly exposed to, and try out, cutting-edge technologies. The study results, as well as the extensive studies of non-expert users in the literature (e.g., [39, 57, 59]), served as fodder for our design work and highlighted opportunities that exist when different subsets of users employ physical and digital tools. Further, because participants' work focuses on innovation in interaction, their own reflections on their work patterns during the task served as further inspiration when developing an innovative set of interaction methods.

Task

Participants were asked to solve the VAST 2006 Symposium Contest [11]. Participants' goal was to find the truth behind a strange news story and summarize evidence from 230 news stories, 3 reference documents, 4 images, and 1 spreadsheet. Most of the documents were not relevant to the task, thus simulating traditional workflows where one needs to seek and read literature, manage documents, and synthesize the document information. The documents were provided as physical printouts, as well as in digital forms. Previous work has made use of the same analytical task to study the use of digital tabletops [21], digital tablets [14], and large screen displays [1]. Here, the task was used to understand how users read and manage documents from both digital and paper mediums to inspire the augmentation of physical documents in mixed reality.

Procedure

Demographics and daily computer usage were first collected from participants through a pre-study questionnaire. The study was then conducted in the participant's own working environment, with signage placed around participants' desks to discourage their colleagues from interrupting them during the study. Participants used their own computing devices and could use any other tools they had in their normal working environment. As in [1, 14], they were also provided with a whiteboard, blank paper, and colorful pens. Participants received digital and physical copies of all 238 documents. Participants could choose to receive the digital copies on either a USB drive or their preferred online cloud platform.

Previous experiments using groups of collaborating participants suggested stopping the task after 90 minutes [21, 32]. As it would take longer for single participants to solve this problem, previous studies allotted a maximum of four hours for completion [1, 14]. However, as interaction behaviors, rather than the participant's ability to complete the task, were of greater interest, the present experiment ended after 60 minutes. Observations from an earlier pilot study revealed that participants become familiar with the tasks and materials within an hour, so this duration of time was sufficient to summarize their working behaviors.

At the end of each session, the participant presented their findings and participated in a semi-structured interview for 20 minutes. Participants were asked why they used, or did not use, each artefact on their desk. The participant was then asked whether they physically arranged the artefacts on the desk using a specific schema or metaphor. Participants were also asked participant-specific questions about behaviors that were noted during the activity, e.g., why one participant created a special folder on a dedicated monitor. Lastly, participants were asked general questions about the limitations of their current environment (i.e., hardware and software) and were encouraged to propose ideas to improve their working environment so that such document-intensive tasks would be easier.

Data Collection

During the research study, a member of the research team observed every session and took field notes about the tools that were used, how devices were organized, and so on. These notes were taken from a distance to avoid distracting the participant. These notes were used to inform the semi-structured interview, which probed participants' utilization of computing devices and physical tools during the task. All study sessions were videotaped and a total of 10.67 hours of video was analyzed via open coding. For each participant, one researcher analyzed the video to understand how participants merged the paper and digital documents into their workflow, how they made use of their workspace, which devices were used for different parts of the task, and so on. These observations were then aggregated, and affinity diagramming was used to identify common themes and unique behaviors that manifested. Given the nature and scale of the design study, qualitative data analysis, rather than quantitative statistical analyses are provided.

Observations

All participants had their own L-shaped desk, as well as a laptop or desktop computer and at least a 24-inch monitor on their desk (Figure 2). They all used paper and digital documents during the study and all of them read the three reference documents on paper. As for the 230 news stories, six out of eight participants used their computer most frequently. While taking notes, six participants used paper, P3 used the provided whiteboard, and P8 created a digital document for notes.

When the video data and participants' interview answers were reviewed, concrete themes emerged regarding how different mediums were selected and impacted workflows. The behavior patterns that were observed, in addition to comments from participants during the interviews, largely fell into three categories: reading activities, spatial organization, and enhancing physical documents.

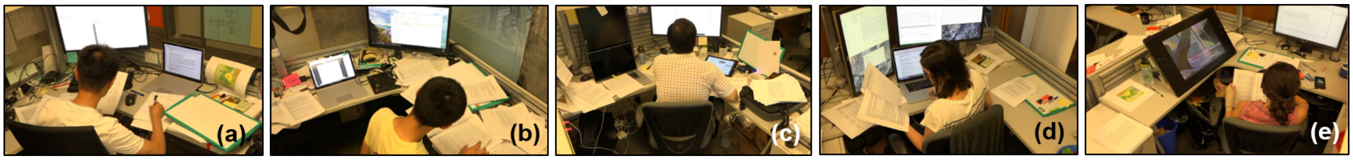

Figure 2: Selected screenshots of participants seated at their L-Shaped desks. (a) P2 taking notes on paper while holding a reference document in his left hand. (b) P4 browsing multiple piles of documents. (c) P5 reading documents on their monitor. (d) P6 using the trackpad while referring to the documents in her left hand. (e) P7 leaning back in her chair and reading documents.

Reading Activities (A). We are interested in the searching and cross-referencing activities during the reading process.

Searching (A1). Six participants used their computer to skim through the news stories. These participants were more likely to use the digital rather than physical versions of the documents due to their ability to search through them, filter irrelevant information, and navigate to target stories. As the paper documents did not enable for quick and rapid document location, the two participants who relied on the printed copies spent much time leafing through piles of documents. These participants, however, noted that paper was easy to navigate, while on the computer they had to *"scroll and open every document and close it"* (P8) and *"computers sometimes got unrelated searching results, so I just manually go through the papers ... [and I am] more familiar with reading text on papers instead of on screen"* (P4). For these participants, the quick visual skimming afforded by paper enabled them to quickly review information in context and utilize the spatial layout of the document to find the desired information. These results suggest that combining the benefits of both mediums might be helpful, because they could provide different functionality while maintaining the flexibility of physical paper.

Cross-referencing (A2). All participants relied on the printed reference documents during the task. P6 commented that *"it is useful to have these pieces of paper just for references to be around ... [I can] quickly check useful references"*. P5 argued that paper was *"flexible to be put anywhere"*. This echoes Bondarenko and Janssen's finding that printing is the easiest way to combine documents from various sources to "one, easily accessible place" [3]. Other techniques, such as the use of physical landmarks on one's desk to aid in organization and retrieval were also observed, e.g., P2 placed the maps and other images vertically at the edge of his desk for reference. The manners in which participants organized and handled the physical documents speaks to the materiality and spatial nature of paper, which can be placed, moved, grouped, hidden, or made visible very quickly. Compared to searching (A1), users preferred to read from paper if they did not need to perform intensive searching or filtering tasks.

Spatial Organization (B). Various patterns were found while participants were organizing physical documents, creating ad-hoc digital storage, and mixing physical and digital documents.

Physical Document Organization (B1). The video revealed that a variety of spatial patterns were used to organize documents in the physical world. For example, five participants grouped the documents into different piles and spread the piles across their desks. Throughout the task, some participants rearranged the piles, placing information that was more relevant closer to their computer, fluidly changing the location and content of piles as the activity continued. Participants sometimes spread the piles over their desk to compare the piles, allowing them to ensure that *"[the documents] are always available"* (P6) and *"[I] can have a glance at everything all at once"* (P7). Although the participant created this structure themselves, some participants had problems finding the correct document or pile. Though searching could be solved by using paper tablets (e.g., PaperTab [58]), interactions for handling the connections between different piles of documents, as well as proper interactions within a pile, remain a challenge.

Ad-hoc Digital Storage (B2). Different strategies to store digital documents were also observed. Five participants created a special digital area to store their discoveries during the study, such as in a certain location of the screen or on a dedicated monitor (if they used more than one monitor). In such spaces, participants created virtual stacks of documents based on theme or content. Unlike the printed documents, participants did not utilize spatiality to indicate importance or recency. To some degree, the partitioning of the screen could be considered to be the creation of working or "storage" zones, echoing prior observations of "storage territories" [50]. Two participants mentioned that they need a larger space to view documents and organize their thoughts, akin to an "analyst's workstation" [1]. Three participants suggested that having a wall-size vertical whiteboard would be helpful when performing such document-intensive tasks in the future.

Mixture of Physical and Digital Documents (B3). One unusual observation involved how the digital and physical copies of a document were organizationally connected. Two participants mirrored the digital organizations they created within their desktop using the printed documents, finding relevant documents that contained important evidence and piling them. Borrowing Hausen et al.'s desktop metaphor [15], these ad-hoc storage spaces served as temporary, secondary workspaces that contained a subset of the artefacts related to the main activities but were located in (virtually or physically) distinct spatial locations. Those who did not utilize spatial piles or collections kept all the documents within the same folder and used searching and a Notepad-style list to keep track of the files they thought were relevant. Since storage spaces facilitate the retrieval of relevant documents, they underline the value of supporting the archiving and retrieving of resources anytime and anywhere.

Enhancing Physical Documents (C). Participants prefer to enhance the physical documents with more functionalities.

Interactive Paper (C1). Although all participants interacted with the printed and digital documents, there was an overwhelming desire for printed documents to exhibit elements of interactivity and not be just "hard paper" [5, 24]. P6, for example, annotated their printed spreadsheet with notes and opened the digital copy on their monitor to create several figures to visualize the contained data. P6 also took notes in a notebook. In the post-study interview, P6 explained that she wanted to augment her notebook to *"link the notes on the paper to the location where I found the evidence"*. P2 wanted a way to *"transfer the notes and paragraphs from the paper to other devices by copy and paste"*.

In addition to annotation, paper was used to assist with wayfinding. P4 made frequent use of sticky notes and used them as bookmarks. He explained, *"It can remind me where I have seen these details"*, but also noted that he *"want[ed] some interactive sticky notes, that [he] can search"*. This desire speaks to the need to maintain the benefits of paper materiality (e.g., for skimming and natural input), while also integrating digital functionality that users have come to value.

4 HOLODOC

A key takeaway from the design study was that combing the affordances of physical paper with the benefits digital technology would improve user experiences (observations A2, B3, C1). Inspired by prior studies of non-expert users (e.g., [39, 57, 59]) and our design study, we designed HoloDoc, a mixed reality system that leverages the advantages of physical artifacts and virtual content, with rich interactions via hand gestures and digital pen strokes. HoloDoc is currently anchored within the context of an academic literature synthesis sandbox. Following Milgram and Kishino's "reality-virtuality continuum" [38], we implemented a prototype to support academic reading with three major techniques that fit into this framework: interactive paper, interactive sticky notes, and an interactive whiteboard.

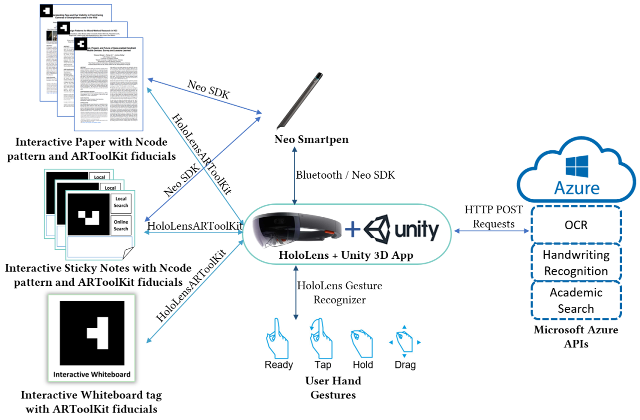

Figure 3: The HoloDoc Architecture. By processing the camera frames in real time (15 fps), HoloLensARToolKit and HoloLens Gesture Recognizer were used to recognize the documents and hand gestures. Pen stroke data was sensed and shared via Bluetooth using Neo SDK. The handwriting recognition, OCR, and academic search functions were supported by Microsoft Azure API via HTTP requests.

Interactive paper is a printed document, which could be used alone in the physical world, while at the same time could be enhanced with virtual content to provide richer functionality. Thus, users no longer need to switch between their digital and analog devices for common activities such as searching (observation A1) and they can access additional and dynamic content quickly and easily (observations A2, C1). The interactive sticky notes serve as a proxy for virtual content in the real world, but still have space for general writing. They rely on feedback in the virtual world so as a result, their physical forms can be manipulated and repositioned as needed, while the associated virtual content remains available (observation B1). The interactive whiteboard, on the other hand, is mostly virtual, enabling for the convenient indexing and grouping of digital content. Users can create a storage space for their digital and even physical content (observations B2, B3). Although academic literature synthesis was explored, the interaction techniques could be generalized to other non-academic activities as well.

HoloDoc Architecture

HoloDoc consists of a Microsoft HoloLens [37], a Neo Smartpen [40] and physical paper documents (Figure 3).

Neo Smartpen-specific (Ncode) patterns were printed in the background of each page of each document using a desktop printer. These patterns, in combination with the Neo SDK, were used to sense users' pen strokes. A fiducial marker was also printed to each document so that the movement

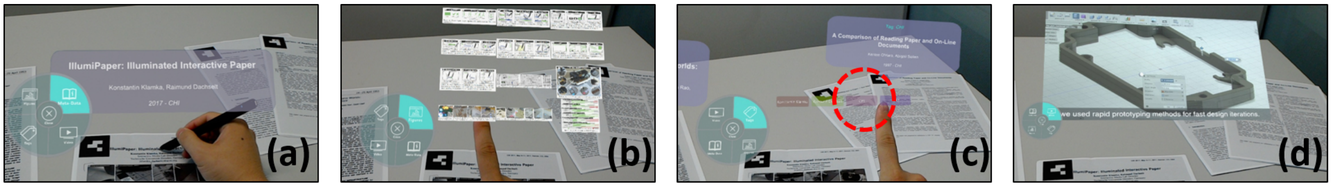

Figure 4: With Interactive Paper, (a) the user taps on a paper's title with the digital pen to invoke a virtual window which contains meta-data information and a pie menu with additional functions. (b) If the user taps the "Figures" button, all the figures and tables from current document are displayed. (c) When the user taps the "CHI" tag, she can identify all documents with the same tag. (d) The user can also play the paper's associated video by tapping the "Video" button. (The red dashes were added here for emphasis - they are not visualized within HoloDoc).

of the document could be tracked throughout the environment, in real time (i.e., 15 fps), by the HoloLens via the HoloLensARToolKit library [46]. Interactive Sticky Notes were also printed with Ncode patterns and fiducial markers to work with said libraries. Interactive Whiteboard tags only have fiducial markers to be tracked by the HoloLens (Figure 3).

Once the pen strokes were recognized, they were relayed to a Unity 3D program running on the HoloLens via Bluetooth. The strokes were then passed to the Microsoft Azure API via HTTP POST requests for handwriting recognition [36]. The Microsoft Azure API was also used for optical character recognition (OCR) in addition to academic search functionality. Users' hand gestures were recognized by the HoloLens' built-in Gesture Recognizer.

To display virtual PDF files and videos, two Unity plug-ins were also used (i.e., the YouTube Video Player [30] to obtain the streamed video online, and PDF Render [43] to download and render each PDF file as a texture in Unity). The current system uses pre-processed screenshots of the webpages of interesting projects and created the links between printed documents and online resources a priori, because rendering a website in Unity 3D and parsing the PDF file were out of the scope of the project.

HoloDoc Interactions

To reduce the abundance of ink traces made on the paper, digital pen strokes within HoloDoc trigger simple actions (e.g., online search) or activate or hide virtual content (e.g., interactive in-text citation menus). When multiple functions are available via these strokes, HoloDoc presents the default option in a virtual window, and hand gestures are used to select the desired option from a pie menu in the bottom-left corner of the field of view. The content of the pie menu depends on the context, i.e., which document the user is working on and which element the user is interacting with. HoloLens' built-in air tapping and drag-and-drop gestures are used to explore virtual content within the environment. After selection, the pie menu rotates the target button towards the window to indicate which content is active. When the user's hand is not detected, the pie menu shrinks to avoid occluding the main content.

Virtual windows can be rendered in various ways in mixed reality and are often affected by design variables such as anchoring-to-world versus anchoring-to-camera and fixed postures versus dynamic postures. Through a few iterations and pilot studies, virtual windows were attached to the physical location and orientation of the element being interacted with so that the user could see the content without distortions.

Interactive Paper

Although physical paper can provide better support for reading activities, due to its flexibility and tangibility [3, 42], according to our design study, it lacks the ability to be indexed and fails to convey rich media (observations A1, A2, C1). To overcome this, HoloDoc enables a user to tap on the title of a document with her pen to explore rich content (Figure 4a). By default, the meta-data of the current document will be displayed, along with a pie menu with additional possible actions: Figures, Tags, and Video. The user can access all the figures and tables inside the document by tapping the "Figures" button (Figure 4b). The "Tags" button visualizes auto-generated tags including the authors of the paper and the publication source (Figure 4c). If the user clicks on a tag, all documents with the same tag will be highlighted by virtual canvases that contain the meta-data information. An enabled "Video" button indicates that this document also has an associated video (Figure 4d).

Because HoloDoc is aware of the location of physical documents, as well as the pen's nib coordinates, additional information is also embedded around the user's focus point. For example, as a user continues to read a document, she might be interested in one of the citations. Instead of flipping to the end of the document and searching for the reference entry, she can tap the pen on the citation to invoke a virtual window that displays the meta-data for that document

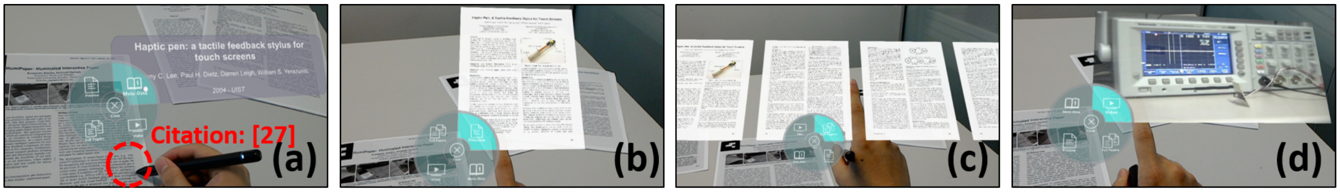

Figure 5: While exploring a reference document in HoloDoc, (a) the user taps on the citation to invoke a virtual window. (b) The user can then tap the "Preview" button from the pie menu and the first page of the cited document will be displayed. (c) Tapping the "Full Papers" button shows the entire document, which can be navigated using hand gestures. (d) The user can also play a video associated to the document. (The red dashes were added here for emphasis - they are not visualized within HoloDoc).

(Figure 5a). The user can also choose to have a quick glance at a document by tapping the "Preview" button (Figure 5b), or spend more time reading the document by tapping on the "Full Papers" button (Figure 5c). Then, she can navigate through the document using hand gestures. The video button is enabled if an associated video is found (Figure 5d). As for other types of references, if the selected reference entry is a webpage, HoloDoc will provide a screenshot of the associated website.

Interactive Sticky Notes

Other form-factors, such as the sticky notes, can also be augmented by mixed reality content. Lepinksi et al., for example, projected contextual interactive menus on sticky notes to provide a ubiquitous way of operating everyday objects [27]. In HoloDoc, a "search window" function is presented via customized sticky notes. When a user wishes to learn about a new concept, she can write down keywords on a sticky note and tap her pen on the "Online Search" button that appears on the note (Figure 6a). HoloDoc will recognize the strokes and present search results in a virtual window that is anchored to the sticky note. The user can interact with the navigation bar to browse the search results or open the associated document in mixed reality (Figure 6b).

These sticky notes are thus a proxy of the virtual search window in the physical world, fitting within the "reality to virtuality continuum". They also helps users build their own knowledge representations via direct manipulation, extending prior results on constructive visualization using tangible tokens [20].

As the user continues to read, she can move the note throughout her environment to reorganize her workspace (observation B1). She can also search from local directories (i.e., already printed documents), by writing down keywords and tapping the "Local Search" button (Figure 6c). A virtual window will then appear, and if a document contains a searched keyword, a virtual indicator will appear above the document, presenting meta-data, such as the title, authors, and publication source (Figure 6d). The indicator's size also varies according to the number of occurrences of the target keyword, thereby enabling the user to easily find the documents which relate most to the target keyword.

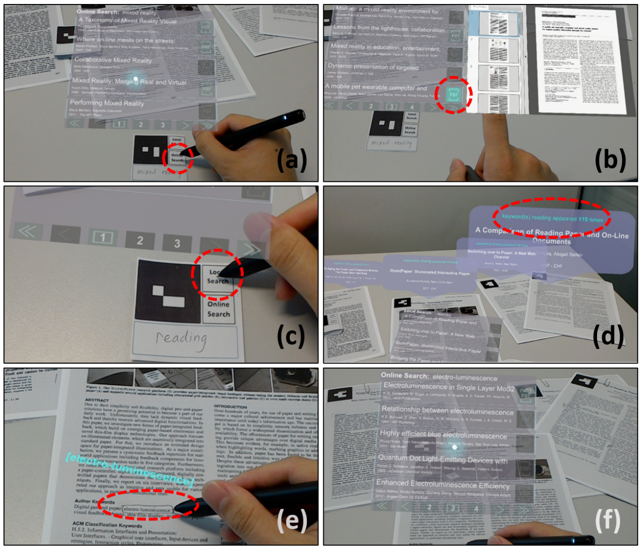

Figure 6: When using an Interactive Sticky Note for searching, (a) the user writes "mixed reality" on the sticky note and taps the "Online Search" button. A virtual window is then displayed and lists the search results, which the user can interact with via the navigation bar at the bottom. (b) The user can also open a relevant document and tap the "PDF" button to browse the document. (c) The user writes "reading" and invokes a local search by tapping on the "Local Search" button. (d) If a document contains the searched keyword(s), a meta-data view will be displayed, varying in size based on the number of occurrences of the keyword(s). (e) If the user circles "electro-luminescence" on the document, (f) she can tap on it again to invoke the online search. (The red dashes were added here for emphasis - they are not visualized within HoloDoc).

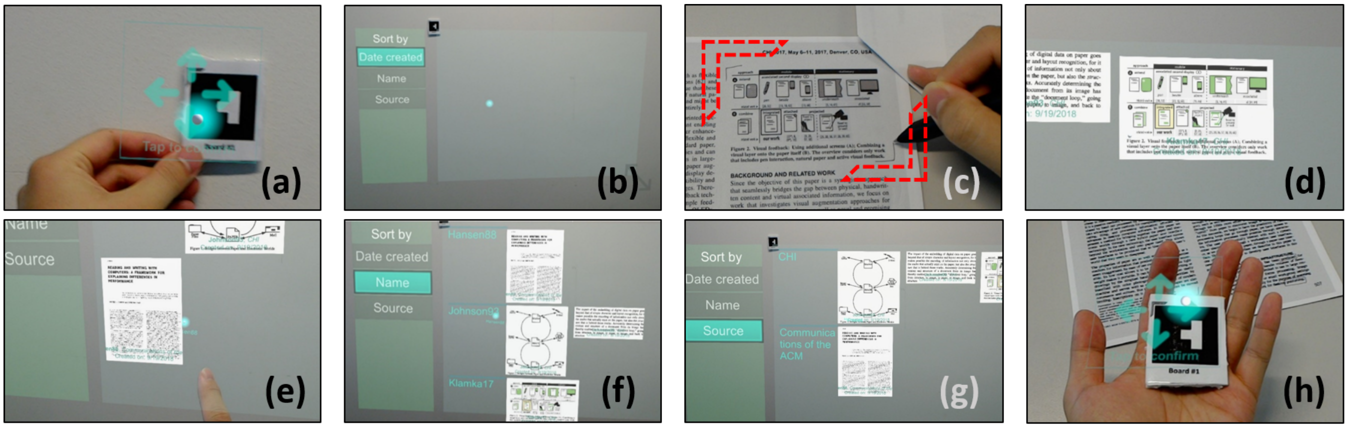

Figure 7: Creating and managing previews from the Interactive Whiteboard. (a) When a user places the whiteboard tag on the wall and taps to confirm its placement, (b) an empty space will be created, with several control buttons on the left. (c) The user can add an interesting paragraph to the whiteboard by cropping it with the digital pen. (d) The content of this paragraph will appear on the virtual whiteboard immediately. (e) If the user finds an interesting reference document, she can tap on the whiteboard to save a preview of it. (f, g) The user can sort the items on the whiteboard in various ways, as shown on the menu. (h) If the user taps on the whiteboard tag to archive its contents, she can bring the tag with her to resume her work anywhere. (The red dashes were added here for emphasis - they are not visualized within HoloDoc).

When the user finds an interesting keyword from the printed document, she can circle the keyword to initiate recognition (Figure 6e), and tap the word again to invoke the online search function (Figure 6f). Other alternative functions could be integrated using this paradigm, such as showing the definition of the selected word(s) or searching for its occurrences in the current document.

Interactive Whiteboard

Based on observation B2, participants wanted a larger space to organize their notes and documents, similar to an "analyst's workstation" [1]. To achieve this in HoloDoc, the user can place a special fiducial tag representing a whiteboard on a nearby wall and tap on it to confirm its location (Figure 7a). This action will create empty whiteboard space (Figure 7b). The user can then use their digital pen to copy and paste any part of a document to the virtual whiteboard (Figure 7c, d). If she finds an interesting reference document while reading, she can tap on the whiteboard while the reference window is open to save a preview of the reference to the whiteboard (Figure 7e). When she has too many previews on the whiteboard, she can organize them by sorting the content (i.e., by creation time, title, or source; Figure 7f, g). If she wants to understand where a preview thumbnail came from, she can tap on the thumbnail to revisit the full document and navigate through it (similar to Figure 5c).

As the whiteboard 'tag' is a proxy for its content in the physical world, similar to the sticky notes, the user can tap on the tag to archive its contents (Figure 7h). Then, when the user arrives at the new location, she can simply place the tag on the wall and tap on it to view all of its collected content.

5 EVALUATION

To understand the utility of an immersive, dynamic tool that mixes physical and digital content, an evaluation was conducted to collect feedback from regular end users, who were researchers from outside Computer Science. Unlike the first study, this population of users regularly performed tasks within the academic research sandbox but were unfamiliar with the technologies in the HoloDoc system. As this population would have a different outlook and frame of reference from the 'experts' in the design study, they were a valuable population to gather first impressions about systems such as HoloDoc.

Participants

Twelve researchers (S1-S12) participated in the study, including 7 PhD students and 5 research-stream Master students (6 female; Mean = 25 years, Range = 22-28 years). Participants came from 10 different departments, all outside Computer Science and Electrical Engineering. All participants represented realistic target users of HoloDoc, as they regularly perform academic reading tasks as part of their daily work. One of the participants had used mixed reality devices (e.g., HoloLens or Glass) before and two of them had experiences using digital pens (e.g., Neo Smartpen or Anoto Pen). Participants were provided with $20 CAD as an honorarium for their time.

Procedure

Before the study, participants were asked to complete a demographic questionnaire about their daily reading behaviors, including the amount of time spent reading from paper, on desktop monitors, and on mobile devices, as well as why they used each medium. Then, the study was conducted in a controlled lab environment. Participants were provided with a HoloLens, a Neo Smartpen, four printed documents, a stack of sticky notes, and an interactive whiteboard tag. All physical documents and sticky notes were tagged with Ncode patterns [40] and ARToolKit fiducials [22].

The study lasted approximately an hour and included two phases. During the first phase, participants were asked to play the role of a graduate student who just started to research how people read documents and were asked to perform a series of research related tasks with HoloDoc (e.g., search for other academic articles online about ‘reading’). The tasks encompassed the three techniques supported by HoloDoc, i.e., the interaction paper, interaction sticky notes, and interaction whiteboard, and all features were explored (Section 4). Participants were able to experience each technique, first through a demonstration by the researcher and then through a free-form exploration process by themselves. During the study, one researcher sat next to the participant to give instructions and answer questions.

In the second phase, a semi-structured interview was conducted. Participants’ comments on each of the techniques were collected to evaluate the usefulness of this system. The interviews were audio recorded for further offline analysis.

Within this study, a reading task, rather than an analytics task similar to the design study, were used because they enabled participants to assess the usefulness of HoloDoc in a familiar context and enabled the session to focus on the tool rather than the user’s competency with the task. As the participants were researchers with varying backgrounds, this enabled them to project and ideate about various elements of the system within their own experiences and use cases.

Results

Answers from the pre-study questionnaire were analyzed to reveal patterns of the participants’ daily reading activities, including the amount of time spent on different reading mediums, as well as the corresponding comments for their choices. A total of 2.42 hours of post-interview audio were analyzed via open coding, with participants’ comments aggregated to identify common themes. Qualitative analyses of the interview are provided.

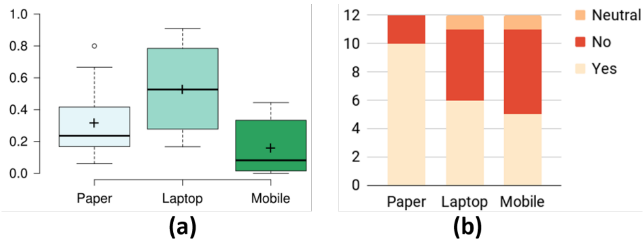

(a) (b)

Figure 8: (a) Distribution of the normalized amount of time spent reading with paper, a laptop, and a mobile device. (b) Participants’ preferences towards reading on these devices.

Reading Behavior. From the pre-study questionnaire, it was apparent that participants read more often on their laptop than from paper or mobile devices (Figure 8a). Interestingly, participants still preferred reading on paper (Figure 8b), due to paper being easier for them to take notes or annotate, paper being more comfortable to read, and paper’s portability. S10 also reported remembering things better while reading from paper. As for disadvantages, S11 found it hard to locate information he/she needed in a short time.

These comments echo findings from the literature, wherein users achieved better performance while reading on physical paper [57] and users prefer to navigate and annotate on paper [3, 42]. On the other hand, their feedback also complemented our design study results from ‘expert’ users, supporting our motivation to combine the benefits of digital and physical documents and functionality.

Usefulness. In general, participants provided positive feedback about HoloDoc and proposed suggestions to improve the system or new functions based on their personal workflows. Here we present several themes summarized from the interview.

Interactive Paper. Although most participants had never used the HoloLens or Neo Smartpen before the study, they quickly grasped the interaction paradigm of triggering actions with the digital pen and viewing multimedia information from the headset. Eleven participants mentioned that HoloDoc’s most useful feature was the ability to display reference information by tapping on the corresponding citation on the physical paper. In this way, users *“don’t need to go all the way back to the end, and you will know what it is talking about”* (S2). Participants noticed that similar functions exist in the desktop reading software; however, they were excited to have them with physical paper (S7). Participants also found the extra information that was accessible via the pie menu (e.g., figures, videos, and full papers) to be helpful, as S11 commented that *“pictures are sometimes more useful than text”.* S5 and S9 wanted to select the author’s name with the pen to obtain more papers written by the authors of the paper. Participants also proposed ideas “borrowed” from existing desktop software, for example, viewing other papers recommended by the system, similar to EndNote or Mendeley.

Interactive Sticky Notes. Six participants considered the online search functions with the sticky notes to be very useful, i.e., *"when some piece of word came into mind, you just write it and search it online immediately"* (S6). Participants also proposed several ideas to improve the searching function such as supporting more search engines (S7, S11, S12), searching by title/author/keywords/content (S1, S9), and displaying the abstract on the search result page (S1, S3). Inspired by the current design, where the system displays the corresponding document next to the search result if the user taps on the PDF button (Figure 6b), S3 suggested that the system could also show the abstract. S8 and S10 also wanted to save some of the papers from the search results for later reference or print them on demand.

Interactive Whiteboard. Six participants appreciated that the interactive whiteboard allowed them to store important pieces from an article rather than storing the entire document. Two were also impressed by the ability to save the whiteboard and carry it around. Other than the existing sorting methods on the whiteboard (Figure 7f, g), S11 wanted to sort the content by type (e.g., figure, table, or paragraph) whereas S2 and S8 wished to manually assign topics to their cropped content and sort them using customized topics. Instead of adding more sorting options, S1 and S12 wanted to drag the contents freely on the whiteboard using in-air gestures. S7 and S12 also wanted to draw on the whiteboard and archive their strokes.

The paradigm of storing an entire whiteboard on a physical tag is similar to the interactive sticky notes in that the tag serves as a proxy in the real world. Surprisingly, S8 commented, *"If it (whiteboard) stores to a Word document, I may think my stuff is there, I know it won't be lost, but when it is stored here (the tag), it seems not tangible to me"*. The physical tag usually seems to be more 'tangible', however, for S8, who was more familiar with 'desktop OS', she preferred to have the content exported to a document. S8 also mentioned, *"What I want most is keeping an electronic copy of what I did physically, so I don't need to do it again"*. Similarly, S10 wished to export the whiteboard content to a cloud drive or a USB key. Inspired by their thoughts, future iterations of HoloDoc could display a file hierarchy interface to increase feelings of 'reality' and 'trustworthiness'.

Navigation and Manipulation. After dragging documents around (Figure 5c), six participants suggested making the other elements in HoloDoc more 'interactive' by tilting (S10), dragging (S1, S11, S12), or enlarging/shrinking (S1, S4, S9, S10, S12) them. Some participants were surprised when they found the HoloLens display to be static in the environment instead of following their view, because they wanted to lean forwards to see the text more clearly. Others mentioned they still wanted to enlarge or shrink the content instead of moving closer or further, similar to what they usually did with touchscreens (S1). Though participants wished to have free control over the elements in HoloDoc, it would be challenging to adjust them to their desired location with simple hand gestures. Using environmental cameras, HoloDoc could be updated to create planes and meshes for easier alignment.

6 DISCUSSION

Based on the findings from design and evaluation study, a number of challenges need to be addressed for optimal mixed reality experiences.

Virtual Element Placement in Mixed Reality

According to the evaluation study, participants could easily understand the virtual elements that were shown in their visual field and were able to propose other helpful features based on their personal workflows. They showed a clear preference of manipulating the elements in their virtual world, including moving, tilting, and resizing the objects. Prior work demonstrated that users could leverage spatial memory to retrieve tasks while managing windows in 3D space [26], however, it is challenging to place virtual elements at the right position and rotation for an optimized reading experience. Sorting and filtering features (like those within the Interactive Whiteboard) and other predefined spatial arrangements may help to reduce "digital messes" and the user's workload.

Many other positioning and manipulation techniques could also be possible. With bimanual interaction, users could adjust position and rotation at the same time, as if they were holding a physical cube. Virtual elements could also be attached to physical objects using a "magnet" metaphor. For example, virtual windows could be attached to static surfaces such as walls and desks, providing "always-available" information at fixed positions. In addition, as physical objects such as phones and cups can be moved, similar to the interactive sticky notes, a physical object could be used as a proxy for a corresponding virtual element, enabling easier manipulation. As users have diverse preferences of repositioning virtual elements, and they may change their desires based on the current task, it is important to allow for easy customizations of virtual layouts.

Text Input in Mixed Reality

The built-in soft keyboard was not used in HoloDoc, because it required the user to turn to each character and perform a tap gesture (or press the Bluetooth clicker). A Bluetooth keyboard could be a potential compromise. The handwriting recognition could also be improved. As suggested by

S7, a printed keyboard with the digital pen could be sufficient, however the system would need to provide real-time feedback in case the pen signal dropped.

System Design Implications

As the weight of the headset affects the user's posture, the visual design of the system could also be affected. As physical paper can be held in any posture, participants preferred to lean backward in their seats and read virtual documents by looking upward to relieve the strain on their neck. The limited field of view (FOV) of the HoloDoc also worsens this problem, as participants had to move their head more often rather than just moving their eyes.

In addition, the FOV also affected text readability. In general, HoloLens has a great resolution for displaying text and graphics, however the location and size of text still needs to be carefully designed. In HoloDoc, the user can drag content closer to read it more clearly, but due to the limited FOV, the user must move her head more often to compensate for this. In addition, users often get motion sickness when the text is too close to their eyes (based on pilot study findings). To reduce motion sickness, HoloDoc initially tries to tilt the virtual window facing towards the user but the window does not move afterwards. Recent breakthroughs in the AR/MR industry (e.g., Meta 2 and Magic Leap One), are working to provide lighter devices with larger FOV, which may mitigate some of these challenges.

Object tracking accuracy is also limited by the resolution of the built-in camera, with only a subset of documents being visible at a time. High-resolution cameras could be used to solve this problem; however, this introduces a tradeoff between harnessing the convenience of a standalone mobile headset and the robustness of relying on external sensors in the environment.

The challenges of making a HoloDoc product also involve the availability of academic resources. An optimal implementation would include cooperation from publishers, but it could be possible to build a modified print driver using a public API (e.g., Semantic Scholar API [53]) and the Neo SDK. This would enable HoloDoc to automatically obtain online resources and link them with the embedded patterns on the printed documents.

7 CONCLUSION

In this work, we conducted a document analytic design study to probe how users employ physical and digital documents when both are available. The study revealed the trend of fusing these resources into a tightly coupled workflow that harnessed the benefits of both mediums. Inspired by the takeaways from the design study, the HoloDoc system was created, i.e., a new mixed reality system for academic reading tasks that leverages the benefits of digital technologies while preserving the advantages of using physical documents. HoloDoc contains several features that span the "reality to virtuality continuum": augmenting regular physical artifacts with dynamic functions, attaching virtual elements to physical proxies, and managing virtual content that was transferred from the physical world. An evaluation of HoloDoc was conducted and demonstrated that users who were not familiar with mixed reality devices could easily understand HoloDoc interactions and were able to reflect on their personal experiences to improve HoloDoc. Design implications and challenges that need to be addressed for future work were also discussed.

ACKNOWLEDGMENTS

We thank Dr. Seongkook Heo and Mingming Fan for providing thoughtful feedback, and the members of DGP lab for their valuable suggestions.

REFERENCES

[1] Christopher Andrews, Alex Endert, and Chris North. 2010. Space to Think: Large, High-Resolution Displays for Sensemaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. 55–64. https://doi.org/10.1145/1753326.1753336

[2] Mark Billinghurst, Hirokazu Kato, and Ivan Poupyrev. 2001. The MagicBook - Moving seamlessly between reality and virtuality. *IEEE Computer Graphics and Applications* 21, 3 (2001), 6–8. https://doi.org/10.1109/38.920621

[3] Olha Bondarenko and Ruud. Janssen. 2005. Documents at Hand: Learning from Paper to Improve Digital Technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*. 121–130. https://doi.org/10.1.1.137.7259

[4] Peter Brandl, Christoph Richter, and Michael Haller. 2010. NiCEBook - Supporting Natural Note Taking. In *Proceedings of the 28th International Conference on Human Factors in Computing Systems (CHI '10)*. 599–608. https://doi.org/10.1145/1753326.1753417

[5] Nicholas Chen, Francois Guimbretiere, and Abigail Sellen. 2012. Designing a Multi-Slate Reading Environment to Support Active Reading Activities. *ACM Transactions on Computer-Human Interaction* 19, 3 (2012), 18. https://doi.org/10.1145/2362364.2362366

[6] Raimund Dachselt and Sarmad AL-Saiegh. 2011. Interacting with Printed Books Using Digital Pens and Smart Mobile Projection. In *Proceedings of the Workshop on Mobile and Personal Projection (MP2) at CHI 2011*. http://isgwww.cs.ovgu.de/uise/Forschung/Publikationen/2011-CHI-ProjectiveAugmentedBooks.pdf

[7] David Dearman and Jeffrey S Pierce. 2008. "It's on my other Computer!": Computing with Multiple Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. 767–776. https://doi.org/10.1145/1357054.1357177

[8] Andrew Dillon. 1992. Reading from paper versus screens : a critical review of the empirical literature. *Ergonomics* 35, 10 (1992), 1297–1326. https://doi.org/10.1080/00140139208967394

[9] Katherine M Everitt, Meredith Ringel Morris, A J Bernheim Brush, and Andrew D Wilson. 2008. DocuDesk : An Interactive Surface for Creating and Rehydrating Many-to-Many Linkages among Paper and Digital Documents. In *IEEE International Workshop on Horizontal Interactive Human Computer Systems (TABLETOP '08)*. 25–28. https://doi.org/10.1109/TABLETOP.2008.4660179

[10] Raphaël Grasset, Andreas Dünser, Mark Billinghurst, and Hartmut Seichter. 2007. The Mixed Reality Book : A New Multimedia Reading Experience. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI EA '07)*. 1953–1958. https://doi.org/10.1145/1240866.1240931

[11] Georges Grinstein, Theresa O'Connell, Sharon Laskowski, Catherine Plaisant, Jean Scholtz, and Mark Whiting. 2006. VAST 2006 Contest - A Tale of Alderwood. In *IEEE Symposium on Visual Analytics Science and Technology*. 215–216. https://doi.org/10.1109/VAST.2006.261420

[12] Tovi Grossman, Fanny Chevalier, and Rubaiat Habib Kazi. 2015. Your Paper is Dead!: Bringing Life to Research Articles with Animated Figures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI EA '15)*. 461–475. https://doi.org/10.1145/2702613.2732501

[13] François Guimbretière. 2003. Paper Augmented Digital Documents. In *Proceedings of the ACM Symposium on User Interface Software & Technology (UIST '03)*. 51–60.

[14] Peter Hamilton and Daniel J. Wigdor. 2014. Conductor: Enabling and Understanding Cross-Device Interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. 2773–2782. https://doi.org/10.1145/2556288.2557170

[15] Doris Hausen, Sebastian Boring, and Saul Greenberg. 2013. The Unadorned Desk: Exploiting the Physical Space around a Display as an Input Canvas. *IFIP Conference on Human-Computer Interaction. Springer Berlin Heidelberg, 2013* 8117 (2013), 140–158. https://doi.org/10.1007/978-3-642-40483-2_10

[16] Juan David Hincapié-Ramos, Sophie Roscher, Wolfgang Büschel, Ulrike Kister, Raimund Dachselt, and Pourang Irani. 2014. cAR: Contact Augmented Reality with Transparent-Display Mobile Devices. In *Proceedings of the International Symposium on Pervasive Displays ïijĹPerDis '14)*. 80–85. https://doi.org/10.1145/2611009.2611014

[17] Ken Hinckley, Morgan Dixon, Raman Sarin, Francois Guimbretiere, and Ravin Balakrishnan. 2009. Codex: A Dual Screen Tablet Computer. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. 1933–1942. https://doi.org/10.1145/1518701.1518996

[18] David Holman, Roel Vertegaal, Mark Altosaar, Nikolaus Troje, and Derek Johns. 2005. PaperWindows: Interaction Techniques for Digital Paper. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*. 591–599. https://doi.org/10.1145/1054972.1055054

[19] Matthew Hong, Anne Marie Piper, Nadir Weibel, Simon Olberding, and James Hollan. 2012. Microanalysis of Active Reading Behavior to Inform Design of Interactive Desktop Workspaces. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS '12)*. 215–224. https://doi.org/10.1145/2396636.2396670

[20] Samuel Huron, Yvonne Jansen, and Sheelagh Carpendale. 2014. Constructing Visual Representations: Investigating the Use of Tangible Tokens. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2102–2111. https://doi.org/10.1109/TVCG.2014.2346292

[21] Petra Isenberg, Danyel Fisher, Sharoda A. Paul, Meredith Ringel Morris, Kori Inkpen, and Mary Czerwinski. 2012. Co-located Collaborative Visual Analytics Around a Tabletop Display. *IEEE Transactions on Visualization and Computer Graphics* 18, 5 (2012), 689–702. https://doi.org/10.1109/TVCG.2011.287

[22] Hirokazu Kato and Mark Billinghurst. 1999. Marker Tracking and HMD Calibration for a Video-based Augmented Reality Conferencing System. In *Proceedings 2nd IEEE and ACM International Workshop on Augmented Reality (IWAR'99)*. 85–94. https://doi.org/10.1109/IWAR.1999.803809

[23] Jeongyun Kim, Jonghoon Seo, and Tack-Don Han. 2014. AR Lamp : Interactions on Projection-based Augmented Reality for Interactive Learning. In *Proceedings of the International Conference on Intelligent User Interfaces (IUI '14)*. 353–358.

[24] Konstantin Klamka and Raimund Dachselt. 2017. IllumiPaper: Illuminated Interactive Paper. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '17)*. 5605–5618. https://doi.org/10.1145/3025453.3025525

[25] Hideki Koike, Yoichi Sato, and Yoshinori Kobayashi. 2001. Integrating Paper and Digital Information on EnhancedDesk: A Method for Realtime Finger Tracking on an Augmented Desk System. *ACM Transactions on Computer-Human Interaction* 8, 4 (2001), 307–322. https://doi.org/10.1145/504704.504706

[26] Jinha Lee, Alex Olwal, Hiroshi Ishii, and Cati Boulanger. 2013. SpaceTop: Integrating 2D and Spatial 3D Interactions in a See-through Desktop Environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. 189–192. https://doi.org/10.1145/2470654.2470680

[27] Julian Lepinksi, Eric Akaoka, and Roel Vertegaal. 2009. Context Menus for the Real World: The Stick-Anywhere Computer. In *Extended Abstracts of the 2009 CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/1520340.1520511

[28] Chunyuan Liao, François Guimbretière, Ken Hinckley, and Jim Hollan. 2008. PapierCraft : A Gesture-Based Command System for Interactive Paper. *ACM Transactions on Computer-Human Interaction* 14, 4 (2008), 1–31.

[29] Chunyuan Liao, Qiong Liu, Bee Liew, and Lynn Wilcox. 2010. PACER : Fine-grained Interactive Paper via Camera-touch Hybrid Gestures on a Cell Phone. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. 2441–2450.

[30] Light Shaft. 2018. YouTube Video Player + YouTube API V3. https://assetstore.unity.com/packages/tools/video/youtube-video-player-youtube-api-v3-29704

[31] Paul Luff, Guy Adams, Wolfgang Bock, Adam Drazin, David Frohlich, Christian Heath, Peter Herdman, Heather King, Nadja Linketscher, Rachel Murphy, Moira Norrie, Abigail Sellen, Beat Signer, Ella Tallyn, and Emil Zeller. 2007. Augmented Paper : Developing Relationships Between Digital Content and Paper Augmented Paper : Developing Relationships. In *The Disappearing Computer*. Springer, 275–297. https://doi.org/10.1007/978-3-540-72727-9

[32] Narges Mahyar and Melanie Tory. 2014. Supporting Communication and Coordination in Collaborative Sensemaking. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 1633–1642. https://doi.org/10.1109/TVCG.2014.2346573

[33] Thomas W. Malone. 1983. How Do People Organize Their Desks?: Implications for the Design of Office Information Systems. *ACM Transactions on Information Systems* 1, 1 (1983), 99–112. https://doi.org/10.1145/357423.357430

[34] Fabrice Matulic and Moira C. Norrie. 2013. Pen and Touch Gestural Environment for Document Editing on Interactive Tabletops. In *Proceedings of the 2013 ACM International Conference on Interactive Tabletops and Surfaces (ITS '13)*. 41–50. https://doi.org/10.1145/2512349.2512802

[35] Fabrice Matulic, Moira C. Norrie, Ihab Al Kabary, and Heiko Schuldt. 2013. Gesture-Supported Document Creation on Pen and Touch Tabletops. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems (CHI EA '13)*. 1191–1196. https://doi.org/10.1145/2468356.2468569

[36] Microsoft. 2018. Microsoft Azure Cloud Computing Platform & Services. https://azure.microsoft.com/

[37] Microsoft. 2018. Microsoft HoloLens. https://www.microsoft.com/en-Us/hololens

[38] P Milgram and F Kishino. 1994. A Taxonomy of Mixed Reality Visual-Displays. *IEICE Transactions on Information and Systems* 77, 12 (1994), 1321–1329.

[39] Meredith Ringel Morris, A. J Bernheim Brush, and Brian R. Meyers. 2007. Reading Revisited: Evaluating the Usability of Digital Display Surfaces for Active Reading Tasks. In *Proceedings of the 2nd Annual IEEE International Workshop on Horizontal Interactive Human-Computer Systems (TABLETOP '07)*. 79–86. https://doi.org/10.1109/TABLETOP.2007.23

[40] Neo Smartpen. 2018. Neo Smartpen - From Paper to Digital, Two Worlds in One Pen. https://www.neosmartpen.com/en/

[41] Moira C Norrie and Beat Signer. 2003. Switching over to Paper : A New Web Channel. In *Proceedings of the International Conference on Web Information Systems Engineering (WISE '03)*. 209–218.

[42] Kenton O'Hara and Abigail Sellen. 1997. A Comparison of Reading Paper and On-Line Documents. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '97)*. 335–342. https://doi.org/10.1145/258549.258787

[43] Paroxe. 2018. PDF Renderer. https://assetstore.unity.com/packages/tools/gui/pdf-renderer-32815

[44] Jennifer Pearson, George Buchanan, and Harold Thimbleby. 2013. Designing for Digital Reading. *Synthesis Lectures on Information Concepts, Retrieval, and Services #29* 5, 4 (2013). https://doi.org/10.1598/RT.64.1.2

[45] Claudio Pinhanez. 2001. The everywhere displays projector: A device to create ubiquitous graphical interfaces. In *Ubicomp 2001: Ubiquitous Computing*. Springer, 315–331.

[46] Long Qian, Ehsan Azimi, Peter Kazanzides, and Nassir Navab. 2017. Comprehensive Tracker Based Display Calibration for Holographic Optical See-Through Head-Mounted Display. arXiv:1703.05834 http://arxiv.org/abs/1703.05834

[47] Yann Riche, Nathalie Henry Riche, Ken Hinckley, Sheri Panabaker, Sarah Fuelling, and Sarah Williams. 2017. As We May Ink?: Learning from Everyday Analog Pen Use to Improve Digital Ink Experiences. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '17)*. 3241–3253. https://doi.org/10.1145/3025453.3025716

[48] Charles Robertson and John Robinson. 1999. Live Paper : Video Augmentation to Simulate Interactive Paper. In *Proceedings of the ACM International Conference on Multimedia*. 167–170.

[49] Stephanie Santosa and Daniel Wigdor. 2013. A Field Study of Multi-Device Workflows in Distributed Workspaces. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (Ubicomp '13)*. 63–72. https://doi.org/10.1145/2493432.2493476

[50] Stacey D. Scott, M. Sheelagh T. Carpendale, and Kori. M. Inkpen. 2004. Territoriality in Collaborative Tabletop Workspaces. In *Proceedings of the Conference On Computer-Supported Cooperative Work (CSCW '04)*. Department of Computer Science, University of Calgary, Calgary, Canada, 294–303. https://doi.org/10.1145/1031607.1031655 arXiv:arXiv:1011.1669v3

[51] Abigail Sellen and Richard Harper. 1997. Paper as an Analytic Resource for the Design of New Technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '97)*. 319–326. https://doi.org/10.1145/258549.258780

[52] Abigail J Sellen and Richard HR Harper. 2003. *The Myth of the Paperless Office*. MIT press, Cambridge, MA, USA.

[53] Semantic Scholar. 2018. Semantic Scholar API. http://api.semanticscholar.org/

[54] Hyunyoung Song, Tovi Grossman, George Fitzmaurice, François Guimbretière, Azam Khan, Ramtin Attar, and Gordon Kurtenbach. 2009. PenLight : Combining a Mobile Projector and a Digital Pen for Dynamic Visual Overlay. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. 143–152.

[55] Hyunyoung Song, Francois Guimbretiere, Tovi Grossman, and George Fitzmaurice. 2010. MouseLight : Bimanual Interactions on Digital Paper Using a Pen and a Spatially-Aware Mobile Projector. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. 2451–2460.

[56] Jürgen Steimle. 2009. Designing pen-and-paper user interfaces for interaction with documents. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction (TEI '09)*. 197–204. https://doi.org/10.1145/1517664.1517707

[57] Kentaro Takano, Hirohito Shibata, and Kengo Omura. 2015. Effects of Paper on Cross-reference Reading for Multiple Documents: Comparison of Reading Performances and Processes between Paper and Computer Displays. In *Proceedings of the Annual Meeting of the Australian Special Interest Group for Computer Human Interaction (OzCHI '15)*. 497–505. https://doi.org/10.1145/2838739.2838745

[58] Ap Tarun, P Wang, and A Girouard. 2013. PaperTab: An Electronic Paper Computer with Multiple Large Flexible Electrophoretic Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI EA '13)*. 3131–3134. https://doi.org/10.1145/2468356.2479628

[59] Craig S Tashman and W Keith Edwards. 2011. Active Reading and Its Discontents: The Situations, Problems and Ideas of Readers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. 2927–2936. https://doi.org/10.1145/1978942.1979376

[60] Pierre Wellner. 1991. The DigitalDesk Calculator: Tangible Manipulation on a Desk Top Display. In *Proceedings of the ACM Symposium on User Interface Software & Technology (UIST '91)*. 27–33. https://doi.org/10.1145/120782.120785

[61] Pierre Wellner. 1993. Interacting with paper on the DigitalDesk. *Commun. ACM* 36, 7 (1993), 87–96. https://doi.org/10.1145/159544.159630

[62] Chih-Sung Andy Wu, Susan J Robinson, and Alexandra Mazalek. 2007. WikiTUI: Leaving Digital Traces in Physical Books. In *Proceedings of the International Conference on Advances in Computer Entertainment Technology (ACE '07)*. 264–265. https://doi.org/10.1145/1255047.1255116

[63] Ding Xu, Ali Momeni, and Eric Brockmeyer. 2015. MagPad : A Near Surface Augmented Reading System for Physical Paper and Smartphone Coupling. In *Adjunct Proceedings of the ACM Symposium on User Interface Software & Technology (UIST '15)*. 103–104.