

Understanding and Supporting Long-Form Text in Mixed Reality

by

Colin A. Brett

Bachelor of Arts, Memorial University of Newfoundland, 2021

A Thesis Submitted in Partial Fulfilment of
the Requirements for the Degree of

Master of Computer Science

In the Graduate Academic Unit of Computer Science

Supervisor: Scott Bateman, PhD, Computer Science

Examining Board: Aaron Tabor, PhD, Computer Science, Chair
Daniel J. Rea, PhD, Computer Science
Matt McGuire, PhD, Education

This thesis is accepted by the
Dean of Graduate Studies

THE UNIVERSITY OF NEW BRUNSWICK

August, 2025

© Colin A. Brett, 2025

Abstract

Mixed reality (MR) technology provides advanced spatial computing and display capabilities, yet its implications for reading and navigating long-form text are not well understood. This thesis investigates, through two controlled studies, how MR can support central aspects of the medium: sustaining deep, focused reading experiences and enabling efficient navigation within long-form documents. The first study examines user experience and flow state when reading narrative text in MR. The second explores MR capabilities for spatially organized navigation in long-form documents compared to conventional approaches. Through these studies, this work demonstrates that current MR headsets can support immersive reading experiences, and that spatial navigation systems significantly reduce task load while improving speed and usability. These findings highlight the importance of designing for human factors and how spatial metaphors can inform the development of more effective and engaging MR reading and navigation systems.

Acknowledgements

Firstly, I have to thank everyone in the HCI Lab and SPECTRAL for your camaraderie, curiosity, and collaboration, in research and beyond.

I also want to give a big thank-you to Dr. Daniel Rea, Dr. Aaron Tabor, Dr. Ernst Kruijff, Dr. André Hinkenjann, and all the others who provided feedback, resources, or encouragement along the way. The generosity of their expertise, time, and insight has been invaluable.

It almost goes without saying, but I owe incalculable thanks to my supervisor, Dr. Scott Bateman, for his steady guidance, sharp thinking, and steadfast support throughout all this. His mentorship shapes both the research and the researcher.

Finally, I couldn't have gotten this far without my friends, family, and partner Jessica—their support nor their patience.

Table of Contents

Abstract	ii
Acknowledgements	iii
Table of Contents	iv
List of Tables	viii
List of Figures	ix
Abbreviations	xi
1 Introduction	1
2 Flow in the Mixed Reality Long-Form Reading Experience	4
2.1 Abstract	4
2.2 Introduction	5
2.3 Related Work	7
2.3.1 Digital Reading	7
2.3.2 Reading in Mixed Reality	8
2.3.3 Flow in Mixed Reality and Reading	9
2.4 Designing a Long-Form Reading Experience for MR	10
2.4.1 A Light Descriptive Framework of Reading Devices	11
2.4.2 System	13

2.5	Study Design & Evaluation	15
2.5.1	Participants	16
2.5.2	Tasks	17
2.5.2.1	Short Reading Task	17
2.5.2.2	Long Reading Task	18
2.5.3	Measures	19
2.5.3.1	Flow and Reading Experience	19
2.5.3.2	Reading Performance	20
2.5.3.3	Task Loading	20
2.5.3.4	Post-Experiment Feedback	21
2.5.4	Statistical Analysis	22
2.6	Results	22
2.6.1	Flow Experience	23
2.6.2	Reading Performance	24
2.6.3	Task Loading	24
2.6.4	Device Rankings & Feedback	26
2.6.5	Flow Correlations	28
2.6.6	Summary of Findings	29
2.7	Discussion	29
2.7.1	On the Framework of Digital Reading Technologies	31
2.7.2	Longer-Form Reading	32
2.7.3	Limitations and Future Work	33
2.7.3.1	Exclusion of Physical Books & Audiobooks	33
2.7.3.2	Participant Pool and Generalizability	33
2.7.4	How Spatial Should Reading Be?	34
2.8	Conclusion	34

3	Spatially Organized Interfaces	
	for Document Navigation	
	in Mixed Reality	36
3.1	Abstract	36
3.2	Introduction	37
3.3	Related Work	41
3.3.1	Document Navigation	41
3.3.2	Spatial Memory	43
3.3.3	Mixed Reality for Instruction and Assembly	44
3.3.4	Immersive Display Spaces & Layout	46
3.4	Designing Spatial Document Navigation	47
3.4.1	Experimental System	49
3.4.1.1	Spatial Compact Condition	49
3.4.1.2	Spatial Extended Condition	51
3.4.1.3	Traditional Condition	51
3.4.2	A Note on Field of View	52
3.5	Study of Spatial MR Document Navigation	53
3.5.1	Reference Document and Assembly Task	54
3.5.2	Procedure	55
3.5.3	Participants	56
3.5.4	Dependent Variables	56
3.5.5	Statistical Analysis	58
3.6	Results	59
3.6.1	Navigation Performance	59
3.6.2	Task Load (TLX)	60
3.6.3	User Experience	61
3.6.3.1	User Experience (UEQ-S)	61

3.6.3.2	Evaluative User Ratings	62
3.6.3.3	Interface Rankings	63
3.6.4	Summary of Findings	64
3.7	Discussion and Future Work	65
3.7.1	Improved performance with spatial layouts	65
3.7.2	Was it really spatial memory?	66
3.7.3	Compact or Extended spatial layouts?	68
3.7.4	Broader applications and future directions	69
3.7.5	Limitations	70
3.8	Conclusion	70
4	Conclusion	72
4.1	Discussion & Reflection	72
4.2	Conclusion	77
	Bibliography	102
	Vita	

List of Tables

2.1 Framework of properties matrix. 13

List of Figures

2.1	The reading view of each device condition: a) eReader; b) LCD tablet; c) MR HMD; d) VR HMD.	15
2.2	Bar charts show the mean ($\pm$ SEM) ratings of the Reading Flow Short Scale, with results of hypothesis testing in the tables.	23
2.3	Bar charts of mean ($\pm$ SEM) performance results by condition. Tables show results of hypothesis testing.	24
2.4	Bar charts show Mean ($\pm$ SEM) NASA TLX ratings by condition. Tables show the results of hypothesis testing.	25
2.5	Bar charts of the count of rankings for each condition. Table shows the results of hypothesis testing.	27
2.6	Correlational analysis of flow with (a) Enjoyment and (b) Task Load subscales (only significant results shown).	28
3.1	The study setup, showing a participant selecting a page in the Spatial Extended document overview.	38
3.2	The Reading Pane and Toolbar in Spatial Extended. Pressing the centre toolbar button activates the overview.	50
3.3	Overview states—compact (left) and extended (right). All thumbnails face the user; the active page is outlined in red, and the hovered target is enlarged.	51
3.4	Mean ($\pm$ SEM) completion time, total interactions, and interaction rate (i.e., interactions/second) by condition.	59

3.5	Mean NASA TLX ratings ($\pm$ SEM) by condition.	61
3.6	UEQ-S results: Mean ratings ($\pm$ SEM) by condition	62
3.7	Evaluative ratings across conditions	63
3.8	Post-experiment ranking by condition.	64

List of Symbols, Nomenclature or Abbreviations

<i>MR</i>	Mixed Reality
<i>AR</i>	Augmented Reality
<i>VR</i>	Virtual Reality
<i>HMD</i>	Head-Mounted Device
<i>SFT</i>	Space-Filling Thumbnails
<i>FOV</i>	Field-of-View
<i>HCI</i>	Human-Computer Interaction

Chapter 1

Introduction

Mixed reality (MR) is a popular class of spatial computing technologies that merge physical and digital environments to enable real-time, interactive experiences [1, 2]. MR includes both augmented reality (AR), where digital elements are overlaid onto the physical world, and virtual reality (VR), where users are immersed in fully digital environments. Modern MR platforms enable users to position, scale, and interact with virtual content in three-dimensional space. These systems typically incorporate spatial anchoring, gesture or pointer-based controls, and user movement, allowing for embodied and situated interaction with digital information [3, 2]. Because of these capabilities, MR has found growing application in fields such as training, simulation, education, design, manufacturing, and entertainment. Across many of these domains, textual content remains a core component, and understanding how MR platforms support or transform the experience of reading is a crucial question for the long-term viability of the technology to support a range of tasks.

Long-form text is a sometimes contentious categorization, with debated thresholds for length, or content. For our purposes, we define it simply as a *book*; an extended collection of text-based information and a central mode of knowledge transmission,

as few mediums have had a more transformative influence than the book [4, 5, 6]. However, while books have endured, they have not remained stagnant, evolving alongside technology—from the printing press enabling mass production and dissemination [7, 8] to digitization in e-books and document files, carrying them into the age of home computing [9, 10, 11]. Despite this continuous evolution and the growing ubiquity of spatial computing, long-form documents remain underexplored in MR. Most existing text-focused MR research has concentrated on short texts or UI text elements rather than continuous reading or comprehensive assistive document navigation [12, 13]. Yet, these long documents introduce unique challenges in MR: decisions in layout, display, interaction, and user comfort all interact in new ways. Importantly, MR systems can suffer from limitations in display fidelity and spatial tracking, raising questions about whether they can support users in focusing on the text itself without the technology becoming a distraction. At the same time, MR may offer affordances not possible in 2D displays, such as unconstrained display space and physical placement, hands-free operation, and its immersive nature. Understanding how MR can support fluid and effective reading experiences is a necessary step toward practical adoption in education, work, and beyond.

This thesis investigates how MR can support long-form document use by examining two broad categories of documents: linear texts, such as novels, that are often rendered in reflowable formats (e.g., EPUB) for continuous reading [14, 15]; and non-linear documents, which use layout-consistent formats (e.g., PDF) to preserve structure for instructional documents or reference material [16]. Each category presents its own question for MR adoption: (1) How can MR systems support sustained, flow-dependent engagement with extended narrative texts? And (2) How can spatial interface design in MR improve navigation and workflow in formatted non-linear documents? Answering these questions required separate controlled user studies to examine distinct but complementary aspects of MR reading.

The first study, presented here as Chapter 2, examines long-form linear reading in MR. We investigate how device characteristics—like physical form and display technology—affect reading experience and engagement. While prior work has explored reading in immersive environments, most has focused on short interactions, brief texts, or novel reading interfaces [17, 18]. This study uniquely focuses on MR’s ability to support traditional long-form reading in extended sessions; revealing complex relationships in physicality and display which shape the users flow and experience. The second study, Chapter 3 of this thesis, targets highly formatted reference documents in MR. As these documents are typically used to complement other activities, participants in this study followed non-linear instructions to complete a physical assembly task. We compare existing MR document reading interfaces with two novel spatial systems, drawing on prior work using spatial document interfaces (e.g., “Space Filling Thumbnails” [19]) and from recent MR experiments showing that spatial placement can reduce task load and improve recall [13], to show the advantages that MR-focused spatially-informed designed can have in performance, task load, and user experience.

Together, these studies contribute to the growing body of HCI research on immersive media by addressing an underexplored but foundational task: reading. We present empirical evidence and analysis that deepen our understanding of reading in MR and propose actionable design guidance for developers building next-generation MR reading systems.

Chapter 2

Flow in the Mixed Reality Long-Form Reading Experience

This chapter is a version of the following working paper:

Colin A Brett, André Hinkenjann, Ernst Kruijff, Scott Bateman. Flow in the Mixed Reality Long-Form Reading Experience. In preparation for submission to the 2026 Conference on Graphics, Visualization & HCI. Waterloo, ON, Canada.

2.1 Abstract

Engagement with long-form narrative text is increasingly shifting toward digital platforms, with devices like eReaders and tablets becoming more common. As mixed reality (MR) technologies gain traction as general-purpose computing platforms, they too are becoming viable mediums for sustained reading. However, the effects of MR on long-form reading experiences remain largely unexplored. This study investigates how MR impacts a reader's flow, experience, and task load compared to established digital reading platforms. We developed a cross-platform reading application and

conducted a within-subjects dual-task study in which participants read the same narrative text across four devices spanning two design axes: Display Type (reflective vs. emissive) and Form Factor (handheld vs. head-mounted). Our findings contribute new insights into how device form factor and display type impact flow, reading immersion and workload—highlighting design considerations for future MR reading applications. We also identify and discuss key usability challenges associated with long-form reading in MR environments.

2.2 Introduction

Long-form reading—broadly defined as reading content such as books or articles of more than a thousand words [20, 21]—demands sustained attention and focus. Through this focused engagement reading can induce a state of “flow.” Flow, or the flow experience, is a concept identified by Csikszentmihalyi [22], best understood as a state of near-complete psychological immersion in an activity, in which an individual becomes optimally focused on a task, often losing track of time and self-awareness [23]. In reading, the experience of flow has been linked to deep immersion in the narrative, heightening both enjoyment and comprehension [24]. Supporting the achievement of flow during long-form reading is therefore a desirable design goal.

With continued interest in supporting many forms of information engagement on mixed reality (MR) headsets, across both productivity and entertainment, there is a growing need to understand how best to support extended reading experiences in this medium. However, most existing research on reading in MR has centred on shorter-form reading, augmenting text with interactive overlays, or comprehension tasks, leaving an understanding of the MR prolonged reading experience—and its relationship with flow—largely unexplored [25, 26, 27].

In this work, we address this gap through a comparative study of long-form reading in mixed reality, evaluated alongside baseline devices commonly used for digital reading. Specifically, we developed a cross-platform reading application that enabled participants to engage with the same continuous narrative text across four devices: an e-ink **eReader**, a high-resolution **LCD** tablet, a fully immersive Virtual Reality headset (**VR HMD**), and a see-through Augmented Reality headset (**AR HMD**). Participants completed a within-subjects study, reading a novel continuously under each condition. We assessed reading performance, cognitive workload, subjective impressions, and flow. By holding content, text parameters, and user interface constant across conditions and varying only the device, our study isolates how each technology’s characteristics influence extended reading experiences.

We find that both the VR HMD and the AR HMD lead to significantly higher experiences of flow. We also find that the AR HMD and the eReader result in significantly lower task load (NASA-TLX) than the VR HMD and the LCD tablet. Interestingly, we also observe that the LCD tablet elicited significantly more errors in reading aloud than the AR HMD. Our findings reveal key opportunities and limitations of MR for long-form reading. Through these results, we clarify MR’s viability for extended reading and offer design insights for future spatial reading experiences. This work provides three main contributions:

- Empirical evidence that MR headsets can support the reader’s experience of flow.
- Design trade-offs, usability challenges and recommendations for future MR reading systems.
- A descriptive framework based on form factor and display type that can guide future study of MR reading devices.

2.3 Related Work

2.3.1 Digital Reading

Reading has long been central to our interactions with digital systems. From early terminal-based character-cell displays to the contemporary high-resolution screens of smartphones, tablets, and dedicated e-readers, screen-based reading has become ubiquitous and continues to grow. According to Pew Research’s 2012 study, individuals who read digital books are more likely to read more, and more frequently [28]; while their 2022 study reported that 75% of Americans read a book in 2022, with 30% reading digital e-books—a 5% increase from just three years earlier [29].

The growth in digital reading has prompted extensive research comparing screen-based reading to traditional print, with a divergence between subjective preference and objective performance suggesting a complex relationship. One study using eye-tracking and EEG found no significant differences in comprehension or cognitive effort between paper books and tablets, despite participants expressing a strong preference for print [30]. Notably, older participants showed shorter fixation durations when reading on tablets. Similar findings have been echoed elsewhere [31, 32], suggesting that subjective preference for print may not align with measurable performance differences. The enduring preference for physical books—often attributed to factors such as age, tactile engagement, and familiarity—persists even when readers acknowledge drawbacks such as physical size and weight, cost, or limited accessibility [33, 34]. In contrast, younger and/or more technologically fluent readers tend to show greater openness to digital formats. These patterns suggest that print preference may be more a product of habit and familiarity than the medium itself. To date, little is known about how readers perceive long-form reading in mixed reality, or how it compares to other digital reading devices.

In their book *Why Digital Displays Cannot Replace Paper*, Shibata and Omura [35] argue there are three primary factors undermining the prevalence of display-based reading: **1)** eye-unfriendly display characteristics (e.g., high contrast, flickering, glare); **2)** operational and physical constraints (e.g., displaying entire pages, page turning, device handling); and **3)** concentration disturbances (e.g., blinking cursors, intrusive icons, distracting menus). While these issues have consistently constrained digital reading for generations, technological advancements, intentional design, and the unique physical affordances of MR devices may be able to bring a about reevaluation of these factors.

2.3.2 Reading in Mixed Reality

The study of long-form reading in MR has received limited attention, though some research has specifically addressed dimensions of the experience. Pianzola et al. found that VR reading, using a 3D book model, can enhance intention to read via a serial mediation of transportation into the story world and affective empathy [36]. Similarly, recent work explored “symbiotic” VR reading experiences that use dynamic thematic environments to complement fictional texts. Although users reported enjoying the experience, their attention frequently shifted between the text and the immersive environment, leading to hybrid narrative interpretations [37]. While these studies centre on new forms of storytelling rather than traditional linear text-based reading, they nonetheless highlight the experiential potential of MR for narrative engagement.

Other research has focused on the technical and ergonomic considerations crucial for sustained reading in MR. Studies have shown that “dark mode” presentation (light text on dark backgrounds) can reduce visual fatigue and enhance legibility in both VR and AR HMDs [38, 39]. Text layout and positioning also matter: paragraph-style

layouts in world-anchored positions outperform RSVP (rapid serial visual presentation) and user-fixed placements in terms of both comprehension and comfort [40, 41]. Optimal parameters identified for longer reading tasks in VR include a display distance of approximately 30.5 cm, text angles between 1.1° and 1.3°, and serif fonts rendered at 18 pt equivalency [42]. Additional work on field of view and curvature found that narrower views (90°) improve reading speed, and slightly curved displays (200r) enhance visual search accuracy compared to flat screens [43].

Further studies on reading in VR have examined caption-based narration in immersive VR [44], AR reading while walking [45], text displayed on 3D surfaces [46], and speed-reading interfaces [47]. While these studies provide important design considerations for interactive MR experiences, they have scarcely addressed how MR devices might support flow and usability in sustained reading.

2.3.3 Flow in Mixed Reality and Reading

Flow refers to a state of deep immersion in an activity, characterized by intense focus, a loss of self-consciousness, and distorted time perception [23]. Flow is closely tied to intrinsic motivation and typically arises during autotelic activities—those pursued for their own sake, rather than for external reward [48]. It emerges when there is a balance between perceived challenge and skill, and has been observed in domains ranging from sports and games to artistic creation and long-distance running [49, 50, 51, 52].

Recently, flow has been explored in MR research. Studies show that core flow preconditions, such as a sense of control and challenge-skill balance, can be achieved in immersive environments like VR [53]. Time distortion, a hallmark of flow, has also been observed among VR users engaged in gameplay [54]. Other studies have examined flow in mobile AR learning contexts [55] and VR sports viewing, where

spectators reported higher flow than with conventional displays [56]. These findings suggest that the immersive properties of MR displays can actively support the induction of flow across different activities.

Reading has long been recognized as a flow-inducing activity. Csikszentmihalyi called it “*perhaps the most often mentioned flow activity*” [57], while others have highlighted its close alignment with intrinsic enjoyment and deep cognitive absorption [58]. Flow during reading has been linked to narrative transportation and mental model construction [59], and is more commonly elicited by fiction than nonfiction—especially when the material is self-selected and personally meaningful [60]. Recognizing reading’s unique cognitive and affective qualities, researchers have developed reading-specific flow measures, such as the Reading Flow Short Scale (RFSS) [58], designed to avoid performance-based framing and better capture immersive reading experiences.

The overlap between reading and flow, combined with MR’s demonstrated capacity to induce flow in other domains, raises an important question: can MR support or enhance the experience of flow during long-form reading? This possibility frames our investigation, which seeks to evaluate whether immersive MR environments can deepen engagement, focus, and enjoyment in sustained reading.

2.4 Designing a Long-Form Reading Experience for MR

To investigate how MR technologies shape long-form reading, we designed a cross-platform reading system to enable controlled comparison across disparate technologies. Our goal was not only to prototype a viable MR reading environment, but also to systematically examine how key device properties influence the reading experi-

ence. This section outlines the conceptual framework behind our device selection, followed by a description of the MR reading environment we developed to support immersive narrative engagement.

2.4.1 A Light Descriptive Framework of Reading Devices

Our primary interest is to better understand the reading experience in mixed reality. To do this, we developed a simple descriptive framework to guide the selection of different devices for comparison and to help interpret our results. The purpose of this framework is to distinguish these digital reading devices by their defining characteristics. In particular, we started by considering a comparison between a VR HMD and an AR HMD. However, even then, there were decisions to make: should the AR headset use a see-through optical display or a video see-through display? We examined the defining characteristics of different technologies, and from this we identified a 2×2 matrix defined by two independent design axes: *Form Factor* and *Display Technology*. See Table 2.1.

The Form Factor dimension has two possible values in our study: handheld and head-mounted. Handheld devices act as our baseline and represent devices commonly associated with digital reading. Head-mounted devices are made possible by displays worn on the head, like mixed reality headsets or smart glasses. The Display Technology dimension also covers two values: emissive and reflective displays. Emissive or (light-emitting) displays—common in devices such as smartphones, laptops, and televisions—employ Liquid Crystal Displays (LCDs) or light-emitting diodes (LEDs) to emit light from the display pixels themselves or from a source behind the display area. Light-*reflective* displays, by contrast, rely on external projected or ambient light to illuminate the display; e-ink displays reflect light in a manner similar to printed paper, while transparent AR displays often redirect light through

optical waveguides, allowing users to see both the digital content and their physical surroundings. While emissive displays are often favoured for their visual acuity, producing vibrant colours and high contrast ratios even in dark conditions, reflective displays have demonstrated the ability to reduce eye strain and improve readability in bright environments and during prolonged use [61, 62].

Handheld Reflective: The handheld, reflective condition for our study was a Kobo Sage (N778), which features an 8-inch E Ink Carta 1200 display with a resolution of 1920×1400 , a pixel density of 300 pixels per inch (ppi), and a weight of only 241 grams—making it the lightest device in our lineup.

Handheld Emissive: The handheld, light-emitting device was represented by an LCD Tablet: an Apple 2nd Generation 11-inch iPad Pro (A2228), with a resolution of 2388×1668 pixels at a density of 264 ppi. It weighs 466 grams, uses an LED backlight, includes an anti-reflective coating, and has a 120 Hz refresh rate, with a standard dynamic range maximum brightness of 600 nits.

Head-mounted Reflective: The head-mounted, light-reflective condition was represented by the Magic Leap 2 (M90AA004), a standalone see-through AR headset running Android. It features dual Liquid Crystal on Silicon (LCOS) projection displays at 1440×1760 resolution and 800 ppi. The Magic Leap 2 uses segmented dimming technology—a unique feature that enables localized black rendering by darkening discrete regions of the display. This capability, absent in comparable AR headsets like the Microsoft HoloLens 2, supports greater contrast and visual clarity, especially important in contrast-heavy scenarios such as reading. The device also offloads many components to an external *puck*, allowing the headset to weigh only 266 grams.

Head-mounted Emissive: The light-emitting, head-mounted condition used a Meta Quest 3 VR headset (S3A) with enclosed dual LCD displays, each with a reso-

lution of 2064×2208 pixels and a pixel density of 1218 ppi. To maintain immersion without distraction, we disabled video pass-through and placed participants in a low-salience 3D environment composed of soft, ambient forms and muted tones. The Quest 3 was the heaviest device included, at 515 grams.

A note about the framework – It is important to note that we do not consider this an exhaustive framework, but rather a useful tool to help guide the selection of reading technologies. In particular, it led us to choose two baselines (the eReader and LCD tablet) as points of comparison with the mixed reality devices. We also acknowledge that the technologies are not perfectly uniform across the framework axes. For instance, both the Meta Quest 3 and Magic Leap 2 are head-mounted displays, but the Quest is much heavier, which could have comfort and usability implications. Similarly, the Meta Quest 3 and the iPad Pro both have emissive displays, but their differing specifications could influence clarity and eye comfort. Still, as we argue in the Discussion, the framework provides a useful structure for interpreting results and identifying directions for future work.

<i>Form Factor</i>	<i>Display Type</i>	
	Reflective	Emissive
Handheld	eReader	LCD Tablet
Head-mounted	AR HMD	VR HMD

Table 2.1: Framework of properties matrix.

2.4.2 System

We developed a single, multi-platform reading application that delivered a consistent experience in software, interface, and text formatting across all devices. This allowed participants to engage with the same narrative text under each condition, minimizing confounds. Due to inherent limitations in the eReader’s display and closed software ecosystem, we were unable to fully implement a customized application for that de-

vice. To address this, we opted to design our system around the eReader’s interface, replicating its layout, font, and typesetting parameters for use in our application. The eReader was the only device that did not run our custom application, but our system was designed to mimic its reading interface as closely as possible. In this way, we maintained consistency across all relevant features (see Figure 2.1).

Our system was developed using web technologies (HTML, CSS, and JavaScript) to run within a web browser. The application has two primary views. First, the library view presents a grid of consistently styled, simplified book covers, with hover-activated tooltips displaying book details such as genre, publication date, and a short synopsis. Second, the reading view, accessed upon book selection, minimizes distractions by displaying only the text without a visible user interface or decoration. Text parameters, including font, text size, line spacing, and visible text area, were dynamically adjusted for each device to mimic the display of the eReader.

Text was rendered using the Georgia typeface, a serif font known for its readability even on small or low-resolution displays and commonly used in e-books [63, 64, 65]. In contrast to traditional print media, we employed an inverse colour scheme, or *dark mode*, with light text on a dark background—adopting near-white and near-black tones sampled from the eReader. This decision was informed by prior research indicating that dark mode enhances comfort, reduces eye strain, and improves readability [38, 66, 67]. When in an HMD condition, the reading pane was fixed approximately 30 cm from the user’s head position, with a subtle angle easing to keep it facing the user as they moved (based on optimal parameters identified in prior research [42]), though it could be repositioned and pinned anywhere the reader preferred.

Both handheld devices could be either held or placed on an adjustable stand, allowing participants to adopt a reading posture they felt most comfortable in. For both HMDs, participants received assistance to ensure a secure and comfortable fit.

Across all four device conditions, participants could navigate the text using either the device’s native input (touch or gesture) or a Bluetooth keyboard modified with a physical overlay (constructed from cardboard and tape) to obscure unnecessary keys. Users were encouraged to choose whichever method they felt most comfortable with, and could switch freely at any time.




Figure 2.1: The reading view of each device condition: a) eReader; b) LCD tablet; c) MR HMD; d) VR HMD.

2.5 Study Design & Evaluation

To investigate how different device characteristics affect the long-form reading experience, we designed a within-subjects experiment in which each participant completed two sequential tasks across four device conditions. The order of conditions was coun-

terbalanced using a Latin square design to mitigate ordering effects. Prior to the study, participants were asked to leave their phones and other potential distractions outside the testing room. They then received an overview of the study and the devices, provided informed consent, and completed a brief demographic questionnaire. All sessions took place in a quiet, well-lit room with only the participant and the experimenter present. Devices were placed on a table and introduced one at a time as participants progressed through the study.

2.5.1 Participants

Sixteen participants took part in our study and were recruited from our university campus and through word-of-mouth. Our participant pool comprised 9 women and 7 men, ranging in age from 23 to 40, with a mean age of 28. We collected information on vision and prior MR use to ensure a range of users and to identify any potential discomfort or learning effects. Eight participants reported having normal vision, while eight reported corrected-to-normal vision, and none reported any known visual impairments. For MR familiarity, two participants had never used an MR device, four had very limited experience, and the remaining ten had used MR devices with some regularity; seven of these reported previously experiencing some degree of VR sickness [68].

Participants were also asked about their reading habits and language abilities. Ten participants spoke English as their first language, while six spoke English as a second language. All self-identified as fluent and proficient in English reading, and all study materials were presented only in English. Nine participants self-identified as *readers*, or someone who considers reading an active hobby, while seven did not. The frequency of reading for personal pleasure ranged from daily to less than once per year.

All participants provided informed consent. The experimental procedure was approved by our university’s research ethics board (file number REB-2023-190), and a gift card was provided as a thank-you for participation.

2.5.2 Tasks

The experiment consisted of two tasks: a short reading task (to measure reading performance) and a long reading task (to measure flow and user experience). The two tasks were designed to be completed sequentially for each condition. Before starting each condition, participants were introduced to the given device and its operation. After the short reading task, a brief break was given to collect measures and allow the participant to rest before proceeding to the long reading task.

2.5.2.1 Short Reading Task

The first task involved a short out-loud reading assessment designed to evaluate participants’ reading performance (i.e., to time and count any reading errors that arose from using the device). Participants were presented with a pre-selected passage of text from a book and instructed to read it aloud in whatever manner was most comfortable to them.

To ensure consistency across all reading passages, we meticulously hand-selected and edited a series of texts to meet specific target criteria. Each passage was extracted from Grimm’s Fairy Tales [69] and adapted to precisely 220 words across 11 sentences. Further edits focused on maintaining comparable difficulty levels across all versions. Specifically, each final passage averaged 1.2 syllables per word, had a Flesch-Kincaid Grade Level of 6.4, and a Flesch Reading Ease Score of 85 [70]. These edits were kept as minor as possible, carefully avoiding any impact on the original meaning,

coherence, or structural integrity of the text, ensuring each passage remained fluid and complete when read aloud.

The selection of passages was balanced across conditions using a mutually orthogonal Latin square design. This ensured that each passage was distributed equally among participants across all four experimental conditions.

2.5.2.2 Long Reading Task

Following the short reading task, participants proceeded to a long reading task in which they were given unrestricted time to browse the system’s library view and select a book. They were instructed to choose any book in the system that they would like to read for the remainder of the study. Participants were encouraged to explore the available options thoroughly and could ask for book recommendations or clarification. This selection process served a dual purpose: to allow participants to familiarize themselves with the device and its interaction mechanisms, and to ensure they chose a book that genuinely interested them for sustained engagement, as interest and enjoyment have been identified as antecedents of flow [71, 72].

Unlike the short passages, the book selection was not constrained by strict difficulty or thematic parameters. Instead, we offered a diverse collection of 20 narrative fiction books, with the primary criterion being sufficient length to support extended reading throughout the study. These books were selected from online user-curated reading forums (e.g., Goodreads). The goal was to provide popular and critically well-received novels, encompassing a range of reading levels and styles—both classic and contemporary—to ensure participants could find options aligned with their individual tastes and preferences. Participants were also asked not to select books they had already read, to avoid any potential bias or familiarity that might affect engagement.

To minimize stress and promote natural engagement and flow, participants were explicitly assured that their comprehension and memory of the book would not be evaluated. They were then instructed to read at their own pace while the experimenter left the room to ensure privacy. After 25 minutes—an interval demonstrated to reliably induce flow [73, 60, 58]—the experimenter returned, and participants completed questionnaires about their experience with the device and the selected book.

2.5.3 Measures

2.5.3.1 Flow and Reading Experience

To capture the reader’s flow and better understand their overall reading experience during the long reading task, we used the Reading Flow Short Scale (RFSS), a reading-specific flow questionnaire designed to sensitively measure flow in the context of reading. The RFSS is an adaptation of the Flow Short Scale (FSS) [74, 58].

The RFSS includes eight items spanning two subscales: *Absorption* and *Smooth Processing*, and omits the FSS items related to *Concern*, which apply only to competitive activities. *Absorption* reflects attentional involvement and engagement in the task, while *Smooth Processing* captures the sense of effortless progress and fluency during activity. Together, these dimensions represent core experiential aspects of flow in non-competitive contexts like reading [75]. Participants responded to each item using a 7-point Likert scale (1 = “strongly disagree”, 7 = “strongly agree”) following each long reading task, allowing comparison across the four device conditions.

2.5.3.2 Reading Performance

During the short reading task, performance was assessed through completion time and reading accuracy, quantified as the number of observed errors. For the purposes of this study, an *error* was defined as the omission or misreading of words, mid-sentence pauses exceeding one second, restarts of sentences or segments, or any disruptions that could plausibly affect reading continuity or comprehension. Minor verbal disfluencies such as stuttering or mispronunciations were not considered errors, as they were unlikely to stem from the reading condition itself. Both time and error counts were recorded live by the experimenter present in the room and later verified using video recordings to ensure accuracy.

No performance data (e.g., timing, behavioural logging, or video) were collected during the long reading task. Participants completed this task alone, with the experimenter monitoring unobtrusively from an adjacent room via two-way glass for safety and consistency. This setup was designed to minimize pressure and foster a more natural, immersive reading environment. Participants were explicitly encouraged to read at their own pace and were told that no comprehension test would follow, allowing them to read without concern for evaluation or performance.

2.5.3.3 Task Loading

The NASA Task Load Index (NASA-TLX), a validated and widely used subjective workload assessment tool, was used to understand the effect of each condition on task load after prolonged use [76]. The TLX captures six distinct dimensions of task load: *Mental Demand*, *Physical Demand*, *Temporal Demand*, *Performance*, *Effort*, and *Frustration*. For each dimension, participants rated their experience on a 21-point scale (0–100, in 5-point increments), with higher scores indicating greater perceived load—except for the *Performance* dimension, where higher values reflect

lower perceived success (i.e., worse performance). To improve clarity, we inverted the *Performance* scores in our reporting so that higher values consistently indicate a greater level of the given dimension.

Our assessment used the "Raw TLX" (unweighted) method, following current practice, as weighting procedures are often omitted for consistency and interoperability across conditions and have been found to be ineffective [77, 78, 79, 80]. We computed the unweighted mean of the subscales to produce an overall workload score; this composite score captures general perceived workload while preserving insight into the relative contributions of each subscale. TLX ratings were collected immediately after the long reading task for each condition, while the experience was still fresh.

2.5.3.4 Post-Experiment Feedback

In addition to comments and questions made during and between trials, participants completed a final questionnaire and interview at the end of the study. The questions focused on subjective impressions of the reading experience and the device used. The questionnaire included a mix of open- and closed-ended items, as well as condition rankings across several experiential criteria. A brief interview followed the questionnaire, allowing participants to elaborate on their experiences and provide qualitative feedback on the devices and reading tasks. This combination of quantitative and qualitative measures aimed to capture a comprehensive understanding of the reading experience across conditions.

2.5.4 Statistical Analysis

All statistical analyses were conducted in R version 4.4.1 [81]. Residual normality for continuous dependent measures was assessed using the Shapiro–Wilk test. When the assumption of normality was met, a repeated-measures ANOVA was used. Greenhouse–Geisser corrections were applied if Mauchly’s test indicated violations of sphericity ($\varepsilon < .75$). Significant main effects were followed by post hoc comparisons using estimated marginal means with Holm–Bonferroni correction.

If residuals were non-normal, sphericity was severely violated ($\varepsilon < .40$), or the data were ordinal (e.g., Likert-type questionnaire items), we used the Friedman test. Post hoc comparisons in these cases used Conover’s all-pairs test with Holm–Bonferroni correction. Additional exploratory correlational analyses used Spearman’s rank correlation for ordinal and non-normally distributed data. All p -values were corrected using the Holm–Bonferroni method within each family.

We report test statistics (F , χ^2 , or ρ), exact p -values, and effect sizes: partial eta-squared (η_p^2) for ANOVAs, Kendall’s W for Friedman tests, and $|\rho|$ for correlations, interpreted as weak (.10), moderate (.30), or strong (.50), following Cohen’s guidelines. The alpha level was set at .05.

2.6 Results

We organize the results of our study by measure and task. We begin with performance analysis from the short reading task, followed by workload, experience, and flow measures from the long reading task, and conclude with overall participant feedback.

When applicable, we employ a consistent reporting style consisting of bar charts

illustrating the central tendency and variability across conditions (mean $\pm$ SEM), with corresponding statistical test results displayed beneath each figure. Where significant main effects were found, pairwise comparisons are annotated above the bars using adjusted significance markers (* $p_{\text{adj}} < .05$, ** $p_{\text{adj}} < .01$, *** $p_{\text{adj}} < .001$).

2.6.1 Flow Experience

We evaluated reading flow using the Reading Flow Short Scale (RFSS) across its three dimensions: *Absorption*, *Smooth Processing*, and a composite *Total* score (Figure 2.2). Significant main effects of device emerged for both Absorption and Total RFSS scores, but not for Smooth Processing.

Post hoc comparisons revealed consistently higher ratings for immersive displays. Both AR ($p_{\text{adj}} < .05$) and VR ($p_{\text{adj}} < .01$) elicited substantially higher Absorption scores than the LCD tablet. The eReader rated slightly higher than the LCD tablet, but this difference was not significant ($p_{\text{adj}} = .053$). The Total RFSS score mirrored this pattern: AR and VR significantly outperformed handheld displays (all $p_{\text{adj}} < .05$), while no differences emerged within form factors.




Figure 2.2: Bar charts show the mean ($\pm$ SEM) ratings of the Reading Flow Short Scale, with results of hypothesis testing in the tables.

2.6.2 Reading Performance




Figure 2.3: Bar charts of mean ($\pm$ SEM) performance results by condition. Tables show results of hypothesis testing.

Analysis of performance during the short aloud reading task found no significant differences in completion time across the four device conditions. However, there was a significant difference in error counts, and post hoc analysis showed that the AR HMD resulted in significantly fewer errors than the LCD tablet ($p_{\text{adj}} = .028$); no other pairwise differences were statistically significant (see Figure 2.3).

These findings suggest that overall decoding and articulation ability remained consistent across devices, with a possible exception that the LCD condition may have introduced challenges that increased the likelihood of reading errors.

2.6.3 Task Loading

Results from the NASA-TLX assessments revealed significant differences in perceived workload across the four conditions (see Figure 2.4). Post hoc tests revealed the following pairwise differences:

Physical Demand: The VR HMD was rated significantly higher than all other conditions: eReader ($p_{\text{adj}} = .003$), AR HMD ($p_{\text{adj}} = .005$), and LCD tablet ($p_{\text{adj}} =$

.014). The LCD tablet was also rated as requiring more physical demand than the eReader ($p = .030$).

Effort: Participants reported significantly higher effort with the LCD tablet compared to both the eReader ($p_{\text{adj}} = .011$) and AR HMD ($p_{\text{adj}} = .023$).

Overall: The LCD condition yielded significantly higher overall task load compared to the AR headset ($p_{\text{adj}} = .004$) and eReader ($p_{\text{adj}} = .007$). VR was also rated significantly higher than AR ($p_{\text{adj}} = .015$).

These results highlight how device characteristics shape perceived workload. First, head-mounted displays showed divergent physical load profiles: emissive VR imposed significantly higher physical demands than reflective AR, suggesting that display technology mediates physical load despite a shared form factor. The LCD tablet, meanwhile, consistently elicited greater perceived effort and overall workload than both reflective displays.

Reflective technologies (both head-mounted and handheld) yielded lower subjective workload overall. AR, in particular, had the lowest mean workload ($M = 20.68$), significantly outperforming LCD and VR. Notably, although VR was the most physically demanding, it still scored lower than the LCD tablet on overall workload ($M = 30.99$ vs. 34.27), suggesting a dissociation between physical and cognitive burdens. These findings point to potential attentional costs associated with emissive displays, even when presented in a familiar form factor (i.e., the tablet).




Figure 2.4: Bar charts show Mean ($\pm$ SEM) NASA TLX ratings by condition. Tables show the results of hypothesis testing.

2.6.4 Device Rankings & Feedback

Participants ranked each condition in terms of *ease of use*, *comfort*, *future use*, and overall *personal preference* (see Figure 2.5). Significant differences were found for **Ease of Use** and **Comfort**. Post hoc comparisons revealed that the VR HMD was rated significantly lower than both the eReader and AR HMD for Ease of Use ($p_{\text{adj}} = .003$ and $.026$), and lower than the eReader for Comfort ($p_{\text{adj}} = .024$).

Reflective displays received the majority of first-place rankings across metrics. Notably, while ease of use and comfort showed significant differences, the more subjective measures of future use and personal preference did not, suggesting a more complex relationship between device traits and user preference.

These nuanced distinctions in user experience were further supported by participants' free-text comments. Nine participants mentioned direct drawbacks of conditions, as P2 shared: *"after [using the AR headset] holding the iPad was so annoying!"* and P12 added: *"the quest was so good until the end, then my eyes began to water"*. General physical barriers among HMDs were also noted, such as by P14: *"I was feeling fatigued or disoriented when taking off the headsets."* and P9 remarked: *"I feel like I was just wearing a helmet"*. For the VR HMD, people particularly noted the discomfort, including the headset being hot and heavy to wear for the 25-minute reading session. A few participants also mentioned that their loss of awareness of the physical environment was an annoyance; for example, when they wanted to shift position to get more comfortable, they would bump the table or feel like they might fall out of their chairs. P13 even mentioned that while reading in VR, they put the keyboard down and could not find it again, and they had been using it for turning the pages. Rather than remove the headset to find it, they opted for the hand gestures, which they found much less usable than the keyboard.

Despite these issues, both head-mounted devices showed situational promise. Several

participants envisioned use cases like reading in bed or while multitasking with the AR HMD. P15 described their desire to read in AR while knitting: *“being able to peek down now and then without issue would be super helpful”*. In summary, participant feedback shows that form factor and display technology can lead to divergent experiences; people place different value on physical requirements, subjective experience, and situational benefits. Importantly, the thoughts on a particular device can be complex; some participants noted both the potential for a device and its usability issues. P13 felt that VR HMD was actually good for reading, but that the lack of comfort and other usability issues reduced its overall value for reading, summarizing with *“[the] quest would be good if it wasn’t so bad”*.

In follow-up discussions, several participants directly voiced excitement for the possibility of improved ergonomics and comfort using the MR headsets. One participant, P11, described how their experience with persistent joint pain has impacted their enjoyment of reading; *“I sit down to read and after only a few pages I’m in pain, and I’m bending out my elbows, and pulling my neck, ..., so I just stop [reading]”*. They went on to say that being able to read while lying down, or in a more comfortable position, would be a significant improvement over traditional reading devices. This sentiment was echoed by P14, who said, *“I just want to read in bed without having to prop myself or my book up on [something]. I can just strap this [AR headset] on and lay back!”*.




Figure 2.5: Bar charts of the count of rankings for each condition. Table shows the results of hypothesis testing.

2.6.5 Flow Correlations

To further explore each device’s relationship with flow and the reading experience, we performed exploratory correlation analyses between participants’ flow, task load, and subjective enjoyment ratings. These factors are believed to be positively conceptually linked with flow and its ties to intrinsic enjoyment, repeated activity engagement, and self-efficacy [58].




Figure 2.6: Correlational analysis of flow with (a) Enjoyment and (b) Task Load subscales (only significant results shown).

Overall, we found a significant positive correlation between participants’ reading flow score (RFSS) and their overall enjoyment of the book, supporting the links between flow and reading enjoyment [22, 24, 60, 82]. When analyzed by device, significant flow-enjoyment correlations emerged for the **AR HMD**, **VR HMD**, and **eReader**, but not for the LCD condition (See Figure 2.6). These results indicate that the strength of the link between flow and overall enjoyment of long-form reading appears to be influenced by both the device’s form factor and display technology, with stronger relationships observed in immersive and reflective reading environments.

Flow–task load relationships showed negative correlations. For the **AR** headset, *mental demand* was strongly negatively associated with flow; for the **eReader**, *temporal demand* showed a significant negative relationship; and in **VR**, *frustration* exhibited a moderate negative correlation. These task load–flow profiles suggest

that different mechanisms may interfere with flow depending on the device: *mental load* in AR, *temporal interruptions* in eReaders, and *frustration* in VR.

2.6.6 Summary of Findings

- MR devices (AR and VR HMDs) elicited significantly higher flow ratings than both the LCD tablet and eReader during long-form reading.
- Task load was rated significantly higher for the LCD tablet and VR HMD compared to both the AR HMD and eReader; the LCD tablet also showed higher task load than the VR HMD.
- Participants made significantly more reading errors on the LCD tablet than with the AR HMD.
- The VR HMD was rated significantly less comfortable than the eReader, with participants frequently citing issues of heat and weight.
- Overall preference favoured the AR HMD and eReader over the VR HMD.

2.7 Discussion

Overall, our findings offer strong evidence that MR devices can effectively support the experience of flow in long-form reading. Both AR and VR headsets yielded significantly higher ratings of absorption and overall flow compared to the handheld devices. This aligns with prior work linking immersive technologies with presence and engagement [83, 36, 84]. While immersion refers to the sensory and perceptual quality of the environment, flow represents a state of optimal psychological engagement. Our results suggest that MR enables deeper engagement with narrative text, not only through immersive presentation but also through affordances such as hands-free use and reclined posture, which may reduce physical effort and support sustained focus.

Among the four conditions, the AR HMD and eReader were the most preferred overall, and both received the lowest workload ratings. These devices were also the lightest in our study, and weight appears to be a strong contributor to comfort, particularly for longer sessions. This may explain why the LCD tablet and VR headset, both weighing nearly twice as much, were rated lower in comfort and imposed higher task load. That said, reading performance remained mostly consistent across conditions, with the exception of the LCD tablet, which resulted in significantly more reading errors.

Interestingly, the VR headset elicited high flow scores despite being rated poorly for comfort and task load. While participants found the headset heavy and warm during the 25-minute session, none withdrew from the task, and several acknowledged its immersive potential. Some participants even indicated that the VR environment enhanced their engagement, suggesting that physical discomfort does not necessarily preclude a strong reading experience. This points to the possibility that future VR hardware, if more ergonomic, could support long-form reading more effectively.

The LCD was rated the lowest on many dimensions and also did not support flow as well as the MR headsets. While the iPad Pro is a technically advanced device, with its backlit screen (which is subject to glare) and its comparatively heavy 466-gram weight, the iPad was both hard on the eyes and uncomfortable to hold. This is also in line with the Quest 3 as the VR HMD condition, which weighs 515 grams and also has a backlit screen. The Kobo Sage eReader and the Magic Leap 2 AR HMD were rated the most comfortable and were roughly half these weights. While not captured by our framework, device weight seems to be an essential consideration in long-form reading, because whether hand-held or head-worn, it impacts the physical positions people can use the devices in and how long they can maintain positions [85, 86, 87].

Finally, the AR HMD seems to hold tremendous promise for long-form reading.

It achieved high flow scores, was rated comfortable, and imposed low task load. Although a few users noted minor eye strain, this was less pronounced than in the VR or LCD conditions. The Magic Leap 2’s high-quality LCoS display—with localized dimming and contrast enhancement—may account for this relative comfort. While possibly not equivalent to e-ink in visual ease, it may offer a more comfortable alternative to backlit screens. Furthermore, the see-through display enabled users to shift positions, remain aware of their environment, and even consider multitasking. Participants described plausible scenarios such as reading in bed or while doing various chores, where hands-free, non-obtrusive access to text could be especially useful.

In the remainder of the discussion, we consider more nuanced results, discuss other insights from our work, and provide implications for future work.

2.7.1 On the Framework of Digital Reading Technologies

In an effort to better characterize an open-space of research [88, 89], we conceptualized an initial framework to help focus our efforts. Our simple descriptive framework began by considering the differences between AR and VR devices and was then extended to include baseline technologies. This resulted in two key axes: display type and form factor. While it likely does not capture all relevant properties of reading technologies, it served as an important basis for selecting the particular headsets (one using reflective display technology and one using emissive technology). Further, we observed patterns related to display technology that spanned the form factor (head-mounted vs. handheld) divide. Specifically, we found that participants considered reflective technologies more comfortable for long-form reading. We believe that future work can build on and extend this simple framework to encompass additional properties of reading devices, in order to better characterize the factors that affect

reading experiences, especially as new technologies, including new display types, form factors, and other characteristics, continue to emerge.

2.7.2 Longer-Form Reading

As previously stated, many studies on MR reading have focused on short-form reading and text interactions [41, 36, 25], while in this work, we have turned the focus toward more time spent with prolonged texts. Beyond focusing on longer-form texts, the duration of study and of reading sessions themselves is an area warranting further investigation. Our long-reading task was limited to 25 minutes—an amount of time adapted from similar studies and considered more than sufficient to allow participants to engage with a text and enter flow [60, 58]. However, real-world reading often involves significantly longer periods and repeated sessions with the same text. While one might consider the results around MR and flow attributable to novelty at first glance, recall that if true, this should be counter to the experience of flow. Flow in reading is defined as complete focus and immersion in the content; in this way, shifting attention to the technology would be counterproductive to flow.

It seems likely that at least some discomforts might lessen with familiarity (or improved hardware), whereas other challenges (like visual fatigue from emissive displays) might accumulate over longer sessions. This would mirror how cybersickness is reduced with more experience [90, 91]. Beyond reading duration alone, it is possible that frequent use may mitigate novelty or unfamiliarity effects present with digital reading. A longer-term or in-the-wild study could shed more light on how these environments contribute to reading experience, behaviour, and the adoption of different reading technologies. Future work should consider a longitudinal study design, allowing participants to adapt to the reading medium over extended periods of time and potentially leading to more meaningful insights on behaviour and

engagement with the text.

2.7.3 Limitations and Future Work

2.7.3.1 Exclusion of Physical Books & Audiobooks

Importantly, our study did not include other book formats—most notably physical books and audiobooks. First, we excluded audiobooks, as they constitute a fundamentally different medium: passive, performative, and non-visual, clashing with our primary interest in the visual and interactive aspects of mixed reality technology. However, with the rise in popularity of audiobooks and podcasts, and given that they share some properties with text-based reading, there is exciting potential for future work to compare—or combine—audio and text experiences using similar tasks and setups.

Second—and perhaps more conspicuously absent—are physical books. While we considered this omission carefully, we ultimately felt they were too dissimilar to offer meaningful comparisons within our framework. Practical constraints in layout, typography, and physical interaction, combined with potential bias toward physical books rooted in tradition and familiarity, led us to exclude them and focus on four conditions we could compare more directly. Nevertheless, now that our work has established that MR devices can better support flow than other digital reading devices, it would be valuable for future research to compare MR with physical books.

2.7.3.2 Participant Pool and Generalizability

Our participant pool was largely recruited from our university campus, and while we attempted to recruit a diverse range of participants, the majority of our participants were students or staff at the university, and many from the computer science depart-

ment. This predisposition or interest in technology may limit the generalizability of our findings to a broader population. Further, this might be particularly relevant because older adults spend substantially more time reading than younger adults (US data) [92, 28]. Future work should consider targeting a more diverse participant pool, including individuals from different backgrounds, to better understand the impact of MR devices on the reading experience across different demographics.

2.7.4 How Spatial Should Reading Be?

In this study, we presented an MR reading experience intentionally designed to parallel traditional reading media, minimizing introduced novelty to enable close comparison across platforms. While this was a deliberate choice, the broader possibilities of reading in mixed reality extend far beyond conventional paradigms—from rethinking reading interfaces and controls, such as pagination and spatial layouts, to inventing new modes of interaction and narrative structure [93].

Having established that MR is not only a viable reading platform but one that may even enhance flow and engagement, we believe this baseline opens the door to more speculative and expressive forms of spatial reading. Future work could embrace the affordances of MR to create novel reading experiences—ones that leverage spatial capabilities not just as a container for text, but as an active component of storytelling and reader engagement. Spatial interfaces for consuming, interacting with, and experiencing text remain underexplored and invite richer design investigations.

2.8 Conclusion

This work examined how people can attain a state of flow—the optimal engagement with text—in long-form narrative reading using digital reading technologies. Using a

within-subjects, cross-condition design, we compared flow, reading performance, and subjective experience across four devices representing distinct points in our framework: form factor (handheld vs. head-mounted) and display technology (reflective vs. emissive). We find that the spatial, visual, and ergonomic affordances of mixed reality headsets better support flow than more common reading technologies. Our findings suggest that while reading ability remained consistent, user experience varied substantially across conditions. Reflective displays consistently elicited lower perceived workload than emissive ones, and the AR headset (Magic Leap 2) in particular struck a compelling balance: offering high narrative absorption compared to handheld options, and greater comfort and flow than the more cumbersome VR headset. These results highlight the complex trade-offs between immersion, usability, and task load in extended MR reading.

Together, these findings offer important guidance for the design of spatial user interfaces for narrative media such as long-form text. They underscore MR’s potential to enrich reading experiences through immersion and presence, while also revealing that physical comfort, interaction modality, and visual ergonomics are critical to sustaining that engagement. We envision future MR reading systems that fully exploit these strengths—hands-free reading, flexible spatial layouts, and narrative immersion—while addressing current limitations in device weight, display strain, and input fatigue. As spatial computing matures, reading may evolve beyond the constraints of flat screens and into environments where stories unfold not only across pages, but across space. Our results suggest that if designed thoughtfully, MR could become not just a viable and enjoyable reading platform, but a uniquely compelling one.

Chapter 3

Spatially Organized Interfaces for Document Navigation in Mixed Reality

This chapter is a version of the following paper:

Colin A Brett, Scott Bateman. Spatially Organized Interfaces for Document Navigation in Mixed Reality. 2025 ACM Symposium on Spatial User Interaction. November 10-11, 2025. Montreal, QC, Canada.

3.1 Abstract

In scenarios leveraging head-mounted mixed reality (MR) displays, digital reference documents commonly support physical tasks such as repairing engines or assembling machinery. Using reference documents often involves frequent navigation between sections and pages, making efficient non-linear revisitation more important than the linear reading of articles or books. Unlike traditional document readers, applica-

tions that use layouts supporting spatial memory have previously been shown to aid non-linear navigation and revisitation. However, spatial layouts for document navigation in MR remain underexplored, particularly in scenarios where attention is divided between digital references and primary physical tasks. We evaluated two spatial interface layouts—a compact and an expanded design—against a traditional document reader interface to support a primary physical task. Both spatial layouts significantly improved performance and user preference over the traditional interface. Though largely comparable overall, the spatial layouts diverged in key areas, indicating trade-offs that may guide interface design decisions based on task demands. This research highlights the value of spatial interfaces in MR, emphasizing their potential to facilitate non-linear navigation and spatial memory formation, and calls for continued exploration of spatial layouts beyond conventional desktop paradigms.

3.2 Introduction

Reference documentation—whether in the form of assembly guides, illustrated process diagrams, reference materials, or maintenance manuals—is extremely common to support a range of physical tasks, including construction [94], assembly [95], and machine maintenance [96]. While traditionally paper-based, mixed reality head-mounted displays (MR-HMDs) have received attention as a means to display digital reference documentation [97, 98], because they allow freedom of movement, flexible placement of the reference information, and the ability to easily access information while having hands free to do work.

Previous work has explored a wide range of novel mixed reality (MR) systems, and compared different techniques and factors to provide guidance across diverse physical scenarios. For instance, Azuma’s 1997 survey detailed the possibilities for aug-



Figure 3.1: The study setup, showing a participant selecting a page in the Spatial Extended document overview.

mented reality (AR) in machinery assembly and repair, pointing to instances from printer maintenance to remote guidance for airplane technicians [99]. Henderson & Feiner evaluated an annotated AR guidance system for performing mechanical maintenance routines inside military vehicles to increase task speed and localization [100], and Tang et al. demonstrated, through a controlled 56-step assembly experiment, that spatially-registered AR significantly reduces cognitive and physical workload compared to traditional documentation or unregistered displays [101]. Studies have also shown that placing tutorial videos at task-relevant physical locations can make accessing content faster and easier during larger multi-step activities [102], and that AR systems designed for construction inspection can reduce task load and error rates [103]. While such examples demonstrate a wide range of innovative approaches to support physical task scenarios, digital versions of traditional paper-based reference documentation remain the most common way procedural guidance is provided com-

mercially. Unlike the interactive 3D models or spatially-anchored videos previously proposed in research, digital reference documents are still the dominant form of guidance material. Their persistence is due to their generality, self-contained nature, and comparatively minimal authoring cost; advantages that most MR-specific systems, which require bespoke content or expert input, do not offer. Despite their advantages, mixed reality interfaces for supporting the use of such documentation have not been widely explored or systematically evaluated.

Using reference documentation tends to differ from the linear page-by-page reading style of articles or books. Prior studies show that effective reference documentation usage involves non-linear navigation, such as jumping between sections and frequently revisiting content [104, 105, 106]. Previous work has demonstrated that traditional interactions and interfaces (e.g., scrollbars and page-wise progression) for non-linear navigation are cumbersome and slower than needed, because they are largely designed for linear reading [107, 108]. In particular, previous work has established that document navigation tools that provide a stable, visual overview of documents called Space-Filling Thumbnail (SFT), can support rapid visual search and the formation of spatial memory, improving performance in non-linear navigation [19]. SFT supports non-linear navigation by displaying every page in a larger document as small thumbnails in a 2D grid. Despite the text in thumbnails most often being illegible, users can visually search by recognizing the unique shape of text and paragraphs on a page, or other visual features such as images and diagrams that are still visible in thumbnails. As people find content they are looking for, they rapidly build up their spatial memory of the location of items, without needing to expend any additional effort. Once locations are learned, information can be accessed even more rapidly because the requirement for visual search is reduced or removed. Together, the rapid visual search and knowledge of specific locations for finding information allow for faster non-linear document navigation.

While the advantages of SFT are well-established [109], there are two important gaps in knowledge around its successful adoption for MR-HMDs. First, while there is good reason to suspect that SFT could be successfully employed on current MR-HMDs, there is little information about how it should be designed for this new display environment. Specifically, previous SFT interfaces in desktop [19] and mobile [110] have been limited to the physical bounds of the display, but MR-HMDs have a much larger potential display space. This raises the question of whether an SFT interface would work best if it is made to be more compact (shrinking all page thumbnails to fit within the viewport of an MR-HMD) or expanded (to use larger thumbnails and allowing them to be displayed more comfortably around the user). On the one hand, compact displays could allow faster visual search because the eyes can be used to rapidly search all pages in a document without the need to slowly move the head to view pages outside of the viewport. On the other hand, increasing the display space could allow for larger, more recognizable page thumbnails, making it easier to find pages and form spatial memory.

The second important gap is that spatial interfaces have often focused on more artificial evaluations that establish whether spatial memory can be formed and specific item locations can be remembered. However, the use case for reference documents on MR-HMDs is to have them support primary physical tasks (such as following instructions to build a model). While SFT interfaces have performed well in focused studies, whether they can perform well while people work on a main task, where they are used more sporadically and in a much less focused way, is unknown.

In this paper, we designed a system that allowed us to compare a traditional, linear document reading application with both a compact and an expanded spatial layout for navigating a reference document (See Figure 3.1). Importantly, our experiment centres around a primary physical LEGO model-building task, which requires

frequent non-linear navigation of the instruction manual. We find that when participants use spatial interfaces for accessing reference documents, the primary physical tasks they work on require significantly less time. While there are few differences between the compact and extended spatial interfaces, participants found that the compact view better helped them orient themselves within the document, while the expanded layout’s greater visual presence was more useful for accessing information. Our paper contributes a first demonstration of the utility of spatially stable overviews for rapidly navigating reference documents in MR-HMDs. Further, we show that benefits exist even while performing a primary physical task, suggesting that spatially stable overviews should provide benefit in real-world use.

3.3 Related Work

3.3.1 Document Navigation

Interfaces that support document navigation have been well researched. By document navigation, we refer to the process and interactions used for moving between pages, sections, or to specific locations within a document. Early original advancements to better support reading documents date to the 1st century, which saw the shift from rolled scrolls to a codex of sequential pages. This new “interface” for reading was a momentous development in document navigation, enabling flexible, rapid, repeatable, and spatially anchored access to content—capabilities impossible in the linear scroll [111]. Human-computer interaction research has explored digital document navigation, with a focus on improving the speed and accuracy of users’ ability to find and navigate to specific points within a document. Studies have found that existing electronic document navigation techniques were both slower than their paper counterparts and less efficient at supporting reading that involved endnotes,

thus requiring repeated revisitation [112].

A variety of techniques have been proposed to improve interaction with digital documents. RSVP (Rapid Serial Visual Presentation), for instance, presents each word, or series of words, sequentially in a fixed location to speed up reading [40]. Scrolling-based approaches have also been refined through scrollbar visualizations [107], rate-based scrolling with automatic zooming [113], and scroll-speed-dependent semantic rendering [114]. Further, Cockburn et al.’s Space-Filling Thumbnails (SFT) [19] introduced a new way of visually navigating documents non-linearly. SFT was shown not only to be faster and preferred by users over other techniques, but also to enhance visual search and spatial memory, despite temporarily occluding the reading view by using all available screen space [109]. Crucially, unlike these techniques constrained by 2D screen metaphors, SFT lends itself naturally to spatial expansion in MR environments. Its emphasis on stable spatial layout and visual memory aligns closely with the affordances of MR. As such, SFT directly inspired the design of our own navigation system.

While scrollbars have long been a standard tool for digital document navigation, several studies have highlighted their limitations, particularly in supporting spatial awareness. O’Hara et al. describe scrollbars as creating a “*dynamic relationship between the location of the information and the reference points, such as the window parameters. This can be contrasted with paper where there is a fixed relationship between text and the pages on which they are displayed*” [115]. This lack of stable spatial mapping inhibits users’ incidental memory for location within the document, which can negatively affect both recall and content re-finding.

Little work exists on document navigation in full 3D environments; however, the introduction of the third dimension into document navigation in traditional 2D displays has been studied. Previous research has explored using a perspective-based 3D ap-

proach to position webpages in a desktop environment to resemble an ‘inclined plane’ (or Data Mountain) [116]. This allowed users to find and retrieve information from these pages faster and more accurately than using a “favourites” mechanism, which, unlike spatial memory-based interfaces, relies on users’ active planning and organization for creation and maintenance. While this study, and others, suggest that spatial memory may be enhanced by an additional dimension, consensus is divided, as more recent research has demonstrated that spatial memory in 2D systems is unaffected by the presence or absence of three-dimensional perspectives [117, 118, 119].

3.3.2 Spatial Memory

Spatial memory is the aspect of our cognition that allows us to store and recall information about the world around us. It involves both our working memory—remembering the position of points and objects around us, and long-term memory—building mental maps of those locations and the physical relationship between them [120]. If we can design user interfaces that organize information such that it supports forming strong spatial memories, we can improve search and navigation speed and, in turn, increase a user’s ability to interact fluidly and efficiently with it [119, 121].

The coding of information into our spatial memory can be allocentric or egocentric [122]. Allocentric memory is dependent on the hippocampal regions of the brain and allows us to remember object location in relation to other external objects. This is particularly useful in navigation tasks, as it allows us to remember the spatial relationships between objects and locations, and can be used to form mental maps of an environment. Egocentric memory, on the other hand, relies on parietal areas of the brain and constructs memories of objects in relation to oneself. This is more useful in object manipulation tasks, as it allows us to remember the spatial relationships between objects and ourselves, and can be used to form mental models of an object’s

position and orientation in space [123, 124]. Previous work has explored the value of leveraging spatial memory in many aspects of user interface (UI) design, with studies showing evidence that spatial memory development in 2D UIs rely strongly on allocentric landmark-based memory structures [125], while the 3D spaces afforded by MR systems require an egocentric reference frame [126]. Although these two systems are distinct, they are not mutually exclusive, and both can be used in tandem to form a complete spatial memory of an environment[127].

Work by Scarr et al. has shown that for users to effectively form and exploit spatial memory in user interfaces, “spatial consistency” in the placement of information is paramount [119]. They emphasize that items must remain in stable, fixed positions for users to learn their locations implicitly. When locations are learned through repeated interaction, the information location is learned without explicit effort. Once established, spatial memory is robust and allows for rapid interaction, without needing to perform slower visual search. Conversely, designs that violate spatial consistency inhibit the formation of spatial memory, forcing users into less efficient interaction strategies. Reference documents (such as instruction manuals) are ideal for interfaces leveraging spatial memory, since their content is typically fixed for long periods of time.

3.3.3 Mixed Reality for Instruction and Assembly

Mixed reality (MR) has been widely explored as a medium for hands-on instructional support, particularly in assembly and maintenance contexts. The hands-free capabilities of modern head-worn MR devices (e.g., Meta Quest 3, Microsoft HoloLens) make them especially well suited for such tasks, having been shown to facilitate support and efficiency in manual assembly [128], and increase confidence for inexperienced users [129]. Hou et al. demonstrated that an AR instruction system using

animated 3D models improved task completion time, reduced errors, and lowered perceived task load compared to traditional paper instructions. They also found a reduced learning curve and enhanced working memory performance with AR-based training [130]. User surveys further highlight positive experiences with MR instruction platforms and suggest their potential to become standard tools for training and maintenance. However, these same surveys note persistent challenges, including headset weight, comfort, and stability issues [131].

Another major focus of MR instruction research has been on providing or improving remote instruction between parties. Researchers have created bidirectional MR telepresence instruction systems, which enable the instructor and instructee to share the same virtual space, with positive reception from users [132]. Other research has implemented systems to allow remote experts to provide spatially-anchored visual instructions to students, which resulted in significantly lower error rates and task completion times compared to traditional instructional material [133]. While remote guidance systems have shown advantages for instruction, their reliance on expert involvement limits their applicability in many scenarios.

Researchers have also begun exploring the integration of traditional document-style instructions into MR environments. Prior work has combined features such as digital hyperlinks, interactive page previews, and sticky notes into paper-inspired MR workflows, which users found intuitive and helpful for navigation [134]. Chen et al. introduced an AR system that segmented captured written instructions into discrete steps and linked them to key objects and locations in the environment. Compared to standard digital readers, their system enabled users to complete stepwise tasks more quickly and with lower perceived workload [135]. These examples highlight the emerging potential of MR to augment document-based guidance without fully replacing it—an approach that remains underexplored but holds promise for scalable

and practical deployment.

3.3.4 Immersive Display Spaces & Layout

When creating an MR application, the layout of digital objects requires special attention, especially when compared to desktop environments. Displaying objects in MR involves many established design decisions relating to properties such as positioning (egocentric: positioned relative to the user; or exocentric: positioned relative to an object or world point), mobility (fixed or movable), proximity to the user, and visibility, among others. These decisions on how to display objects in MR are vital to ensuring usability [136]. While some features are limited by hardware, such as field of view, other aspects must be intentionally designed and have interactions built to support them; for example, the positioning and dynamic sizing of a reading pane in an MR document reader.

One particularly relevant dimension of MR layout is the incorporation of the third dimension. Prior studies have shown that traditionally flat displays (e.g., televisions or large monitors), increasing curvature can improve comfort and immersion as display size grows [137]. Liu et al. investigated how spatial layout affects information perception by studying the use of curved versus flat arrangements for displaying small multiples in immersive environments [138]. Their results showed that fully surrounding layouts were least preferred and did not improve interpretability over flat or semi-curved arrangements. In follow-up work, the authors examined how layout curvature can affect spatial memory. Participants were asked to remember card locations presented in either a flat grid or a fully circular layout [120]. The circular layout, though equidistant in physical space, impaired spatial memory and was described as disorienting due to the increased head movement required. These findings suggest that while curved layouts may feel more immersive, they can hinder task

performance if they overload spatial cognition or increase physical effort. Together, this work provides insights that can be drawn upon for designing spatial layouts to support document navigation.

3.4 Designing Spatial Document Navigation

We developed an MR document reading application for our study that provided three navigation interfaces. The baseline interface, called *Traditional*, is representative of common desktop-influenced implementations found in commercial MR document readers. Traditional interfaces for document navigation (such as that in Adobe Acrobat Reader for the Meta Quest 3) provide well-established and familiar navigation methods, including scrollbars, pagination buttons, and numeric page entry (to jump directly to a specific page). Previous work in document navigation has been based on desktop systems, with display space bound by the monitor [19, 139, 119, 140], contrasting with the unconstrained display space (i.e., all space around the user) available in MR. While spatial memory plays a role even in traditional systems—users rely on recalling consistent positions of menus, buttons, and scrollbars—research shows that explicitly designing for spatial organization and stability can further enhance navigation, revisitation, and performance [116, 140]. We therefore developed two distinct spatial interfaces, both inspired by SFT, to allow us to explore variations of the space filled by thumbnails in a physical environment. Below, we describe design factors we considered might influence spatial memory and visual search, particularly through dimensions of scale, geometry, and visibility:

- **Display Extent and Field of View:** While objects may be placed with few physical constraints in MR, users are only able to view objects within the viewport of a head-worn display. Constraining the layout of pages to that

viewport (and in turn, the user’s field of view) allows for immediate access to all content, facilitating global awareness and minimizing physical effort [141, 142], at the cost of reduced visual size and distinction. In contrast, distributing content beyond the viewport creates spatial groupings that can aid search [143], at the cost of increased physical effort and reduced simultaneous visibility [144].

- **Thumbnail Size and Recognizability:** Thumbnail size trades off visual clarity for compactness. Smaller thumbnails enable fast scanning but reduce detail, which may hinder recognition and legibility [145]. Larger thumbnails improve fidelity and may boost spatial encoding, but take up more space.
- **Structure and Landmarking:** Mapping of spatial memory relies on salient, discriminant features to support navigation and orientation [146, 147]. Compact layouts provide fewer exploitable features to anchor these landmarks due to the simpler arrangement and reduced visual detail, but a smaller presentation allows for less area to map and more fine-grained visual navigation. Conversely, larger layouts can facilitate intentional landmarks through spatial arrangement [125], and enforce natural landmarks within thumbnails by increasing size and clarity, despite their overall larger search area.
- **Physicality and Cognitive Cost:** MR introduces physical demands that traditional UIs do not. Layout size scales embodied interaction through head or body turning, which may enhance spatial encoding through sensorimotor engagement. However, in contexts where navigation supports an external task may introduce interruption, increase task load, or split user attention; requiring examination of these costs against the potential benefits.

We created two designs that reflect these factors, and provide natural points in the larger design space of possible spatial layout: (1) *Spatial Compact*, constrains the display of all pages within the viewport by reducing the size of the thumbnails; (2)

Spatial Extended, extends the display of pages beyond the viewport, leveraging the larger display space to provide larger thumbnails.

3.4.1 Experimental System

Below, we describe the experimental system we designed and built to allow comparison of the Traditional, Spatial Compact, and Spatial Extended interfaces to support document navigation. The system was developed using Unity and the Microsoft MRTK3 Mixed Reality Toolkit for deployment on the Magic Leap 2 see-through augmented reality headset. The system was iteratively tested during development to try and reduce confusion or usability issues that can occur with participants unfamiliar with MR devices. Our system used hand-tracked interactions, rather than the Magic Leap 2 controller.

The base interface provides consistent functionality across all conditions: each includes the same **reading pane**—a movable and linearly scalable rectangle displaying the currently selected page of the document—and a *toolbar*, a context-dependent button group element that lazily follows the user at chest height, but can also be freely moved or “pinned” to a static location. Each condition is described in more detail below (see Figure 3.2).

3.4.1.1 Spatial Compact Condition

In this viewport-constrained condition, the toolbar contains three buttons: *Page Forward*, *Page Back*, and, in the centre, an activation button that hides the reading pane and enters the “overview” state (see Figure 3.3). In this view, the system displays all pages of the document simultaneously, positioned within the available active display area of the device.



Figure 3.2: The Reading Pane and Toolbar in Spatial Extended. Pressing the centre toolbar button activates the overview.

When in the overview state, all pages are rendered as thumbnails, sized to fit within the display area on a plane placed at a distance of approximately 60 centimetres. This condition allows navigation to a desired page through direct selection using either near interaction (i.e., poking the thumbnail) or far interaction (i.e., raycast pointing using the hand with a pinch selection). Selecting a page closes the overview and reverts the interface to the toolbar and reading pane, now displaying the selected page.

To enhance readability and navigation within the overview, an active hover interaction was implemented. This interaction follows the user's raycast from their hand and enlarges the page on hover, bringing it slightly closer to the user. This interaction also serves as selection affirmation. Additionally, a red border is displayed on the thumbnail of the currently displayed page, providing a visual landmark, which has been shown to facilitate spatial learning [148].

3.4.1.2 Spatial Extended Condition

The second spatially organized condition begins with a layout and functionality identical to the Compact condition. Importantly, however, the layout of thumbnails in the overview state extends beyond the size of a single viewport, partially surrounding the user (see Figures 3.3 and 3.1).

The Extended condition’s overview state uses three flat planes: one centred in front of the user, with an additional plane on either side rotated 40°, such that the overview fills ~120° field of view from the user. We determined the exact positioning and size through piloting, and included distinct landmarks at the meeting points between the three panes to aid spatial memory formation.



Figure 3.3: Overview states—compact (left) and extended (right). All thumbnails face the user; the active page is outlined in red, and the hovered target is enlarged.

3.4.1.3 Traditional Condition

Our baseline replicated standard navigation features found in existing MR HMD PDF readers. These included a horizontal scrollbar, direct page number entry, and

stepwise forward and backward pagination buttons, all accessible through the toolbar. The scrollbar supported both near- and far-interaction, and quick navigation to any point in the document by selecting positions along the track. The numeric keypad input allowed participants to directly specify a target page.

3.4.2 A Note on Field of View

Field of View (FOV) is a key factor distinguishing spatial layouts; determined by both device optics and interface design. The Magic Leap 2 FOV, and subsequently the Compact overlay, corresponds to a $45^{\circ}\text{H} \times 55^{\circ}\text{V}$ area centred in the user’s view, comparable across see-through AR headsets and, importantly, to the visual angle often occupied by desktop displays at typical distances [149, 150].

Restricted FOVs have been shown to slow visual search and degrade spatial awareness, whereas larger FOVs can improve spatial presence and navigation (at the cost of larger search areas and potential perceptual/comfort trade-offs) [151, 149, 152, 153]. Consequently, the Compact layout will change if scaled with device FOV: smaller FOV devices would force reduced thumbnails and selection granularity, while wider FOV devices (beyond familiar desktop equivalents) permit larger thumbnails with richer contextual elements, influencing recognition speed and memory encoding [145]. Conversely, the Extended layout is designed for user space and not device display; therefore, differing device FOVs would not impact thumbnail size, though they may affect view obstructed by the overview, required movement, and other perceptual factors.

3.5 Study of Spatial MR Document Navigation

We designed our experiment to focus on assessing the potential utility of an interface designed to leverage spatial memory. The ability to form spatial memory using the type of organization and displays in our experimental system has been well established in previous work [119, 154, 155]. However, previous work has not demonstrated whether a spatial document navigation system could work effectively in MR, and there is little information about how such a system should be designed. As previously described, our experimental system was designed to compare two display alternatives: Spatial Compact and Spatial Extended.

In contrast to earlier studies that isolate spatial memory in controlled tasks, we embedded our system within a primary physical activity to reflect a motivating MR use case of referencing documentation while keeping hands free to perform physical tasks.

Therefore, our study focused on two main research questions: 1) Can a spatial interface improve performance and usability compared to a traditional interface while users are engaged in a physical task? 2) Should the interface be displayed compact (i.e., the headset’s viewport) or extended (i.e., leveraging space around the user)?

We developed a within-subjects study in which participants were asked to construct unique LEGO models as the primary task. Our document system displayed a specially designed LEGO model instruction book, which required frequent referencing and revisitation as participants navigated it using one of the three conditions. The sequence of conditions for each participant and the specific LEGO model used within each condition (three LEGO models were developed) were counterbalanced using a Latin square design.

3.5.1 Reference Document and Assembly Task

We designed the instruction document to follow familiar elements from instruction manuals, assembly guides, and academic reference books. The document had 90 pages, and included multiple model assemblies along with standard front and back matter: cover, title page, table of contents, preamble, and appendix. Its visual style drew inspiration from widely recognized LEGO and IKEA manuals, and each page contained a mixture of textual and pictorial information. Assembly steps were presented as sequential black-line illustrations with a consistent visual perspective. Visuals for instructions were created using BrickLink Studio 2.0 [156], with precise 3D LEGO models that enabled high-quality stepwise diagrams for consistency and accuracy throughout.

LEGO assembly was selected as the primary task because of its accessibility, familiarity, and being an engaging task for participants, as well as its natural stepwise instructions. This choice supported the creation of a document designed for repeated re-visitation and non-linear navigation without introducing excessive learning demands or ambiguity. We achieved this by intentionally decoupling key information across sections and pages—a structure common in real-world manuals.

At the start of each build, participants consulted the table of contents to locate the relevant models. Within that multi-page instruction, new pieces were highlighted and placement indicated with arrows for each step. Required components were identified only by a custom alphanumeric ID (a study-specific “*piece-code*” consisting of a type letter—e.g., A = bricks, B = plates—followed by four arbitrary digits). This system replaced LEGO’s standard Element IDs to avoid prior knowledge effects. Participants then used the piece-code to find the part in an alphabetized index in the appendix revealing a colour image and a location number (1–6; corresponding to labelled cups in front of them). Each cup contained both required and distractor

pieces, evenly distributed to prevent location-based strategies. After retrieving the correct pieces, participants returned to the instructions and continued assembly. This design required the frequent movement between sections of the document to complete each model, prompting reliance on a spatial memory of layout, and frequent visual search of pages.

3.5.2 Procedure

Participants were first introduced to the task procedure and document layout. Each trial began with a demonstration of the system’s interface and interaction techniques, followed by a training phase using a “practice” LEGO model, allowing participants to become familiar with the mechanics and workflow of the condition.

During training, participants were required to complete at least three steps but were permitted additional time to practise freely. At the start of each trial, the instruction manual was reset to the cover page, and participants were informed which model they would be constructing. Errors were recorded throughout the task, but no feedback or penalties were given.

The three experimental LEGO models were matched for difficulty and duration. Each included a comparable number of total and unique pieces, consisted of 9–10 steps, and resulted in a recognizable animal figure. Pilot testing confirmed that no model or instruction set was perceived as significantly easier or harder.

Participants were informed that the experimenter would not assist during task execution, though they were free to speak or think aloud. Any verbalized reasoning, difficulties, or interface-related comments were noted for later analysis. After completing all three conditions—with questionnaires and short breaks in between—participants took part in a semi-structured interview to elaborate on their experience and clarify

observations.

3.5.3 Participants

We recruited 21 participants (12 women, 8 men, 1 undisclosed) via email advertisements and direct invitations. Ages ranged from 18–53 years ($M = 26.5$). All had normal or corrected-to-normal vision, and most were right-handed ($n = 18$). About half ($n = 10$) had no prior experience with MR devices; the remainder reported experience ranging from “*limited*” to “*frequent use*”. Seven participants reported a history of virtual reality or motion sickness, though none experienced symptoms during the study. All participants were familiar with LEGO: eight had not used it since childhood, while 13 described themselves as occasional or active hobbyists. No significant differences in performance were observed with respect to age, gender, MR experience, or LEGO familiarity.

All participants provided informed consent. The experimental procedure was approved by our university’s ethics board (REB-2023-190), and participants received a gift card as compensation.

3.5.4 Dependent Variables

The experiment consisted of a single independent variable (navigation scheme), which was assessed via multiple dependent variables, which we group into the following categories: performance, task load, and experience variables.

Performance— We collected and analysed completion time, total interactions, and interaction rate (defined as the number of interactions per second). Completion time was recorded in seconds, measured from the participant’s press of a *begin* button to an *end* button in each trial. An ‘interaction’ was defined as any button press, page

navigation, slider adjustment, UI element repositioning, or a hover dwell exceeding one second in the overview state. We report both total interactions and interaction rate per trial, rather than a per-page or per-step measure, since the task did not always require accessing a new page for each action. For example, participants could sometimes extract information—such as a forgotten colour—from a page thumbnail without fully navigating to the page.

Task Loading – Following each task, participants’ perceived Task Loading (TL) was measured using the NASA Task Load Index (NASA-TLX)[76] and its six dimensions of workload based on subjective user feedback of Mental Demand, Physical Demand, Temporal Demand, Frustration, Effort, and Performance.

User Experience – User experience was measured in three main ways. First, the short form User Experience Questionnaire (UEQ-S)[157] was administered to capture experience relating to “pragmatic” and “hedonic” qualities.

Second, seven 5-point Likert-type ratings were developed to capture the following:

Distraction: The degree the user is distracted from the task.

Transparency: The extent to which the system becomes less intrusive with continued use.

Reorientation: Support for memory of relative location within the document upon returning to navigation.

Support: The capacity to aid task completion.

Disruption: The frequency with which the user must cognitively shift between the task and navigation.

Interruption: The interference with continuous workflow.

Concentration: The support for maintaining focus on the task.

Finally, after using all three conditions, participants ranked them across a variety of

criteria and provided optional free-text comments. Rankings captured perceptions of each condition’s effect on working speed, efficiency, satisfaction, overall preference, and desire to use the system in future work contexts.

3.5.5 Statistical Analysis

We first tested residual normality for each continuous dependent variable using the Shapiro-Wilk test. When residuals were normally distributed, we conducted repeated-measures ANOVA, applying Greenhouse-Geisser corrections when sphericity was violated ($\epsilon < 0.75$). Post hoc comparisons were performed using estimated marginal means with Holm-Bonferroni correction.

When $\epsilon < 0.4$, residuals violated normality, or in ordinal data (e.g., NASA-TLX, Likert items, rankings), we used non-parametric Friedman tests. Significant Friedman results were followed by pairwise comparisons using Conover’s test with Holm correction.

An exception was made for analysis of UEQ-S results, which, while ordinal, are commonly treated as continuous due to their transformation and reliability. In line with prior work [158, 159, 157], we used the parametric methods for these measures. Effect sizes are reported as partial η^2 for ANOVA results and Kendall’s W for Friedman tests. Greenhouse-Geisser epsilon (ϵ) values are also reported where applicable. All statistical results are reported via paired bar charts and tables. In the text, we only report the pairwise differences for tests where the main effect was found to be significant.

3.6 Results

To support interpretability across diverse measures, we adopted a unified reporting style throughout this section. For each applicable metric, we present a bar chart summarizing the central tendency and variability across conditions (mean $\pm$ SEM), accompanied by relevant statistical test outcomes embedded directly beneath the figure. Where omnibus tests were significant, pairwise comparisons are shown as brackets connecting the bars, annotated by adjusted significance levels (* $p_{\text{adj}} < .05$, ** $p_{\text{adj}} < .01$, *** $p_{\text{adj}} < .001$).

3.6.1 Navigation Performance

We found significant differences in each measure of performance (see Figure 3.4). First, in **completion time**, the Spatial Extended interface yielded the fastest performance, followed by Spatial Compact, and finally Traditional. Pairwise comparisons revealed: **Traditional** was $\approx 27\%$ slower than **Compact** ($p_{\text{adj}} = .005$); **Traditional** was $\approx 54\%$ slower than **Extended** ($p_{\text{adj}} < .001$); No significant difference between **Spatial Conditions** ($p_{\text{adj}} = .137$).




Figure 3.4: Mean ($\pm$ SEM) completion time, total interactions, and interaction rate (i.e., interactions/second) by condition.

Differences were even more pronounced for **interaction count**. The Traditional interface required $2.5\text{--}3.5\times$ more interactions than Compact and Extended, respectively (both $p_{\text{adj}} < .001$). Between the spatially organized conditions, Extended required significantly fewer interactions than Compact ($p_{\text{adj}} = .018$).

Interaction rate (interactions per second) also varied significantly, with Traditional exhibiting the highest rate (both $p_{\text{adj}} < .001$). Although Extended had a slightly lower rate than Compact, the difference was not significant ($p_{\text{adj}} = .308$), suggesting that its speed advantage stemmed from reduced interaction frequency rather than faster execution.

Although we did not collect data on the strategies people used, we did observe that through the tasks, people became quicker and more confident with the spatial interfaces. This suggested that people were learning the locations of the different areas they were looking for. With the traditional interface, we observed participants using a combination of the tools rather than just relying on a single tool for navigation. For example, some participants would rapidly scroll to the area of the page in the document, and then make fine adjustments using the page buttons.

Overall, these results indicate that spatial interfaces — particularly Extended — improve navigation by reducing completion time and requiring fewer interactions and a lower interaction rate.

3.6.2 Task Load (TLX)

Results of the NASA-TLX questionnaire are shown in Figure 3.5. Across all dimensions, the Traditional navigation method resulted in the highest perceived workload. Significant differences were found for *Mental Demand*, *Frustration*, *Effort*, and overall *Task Load* (i.e., the Composite Score).

Post hoc comparisons revealed that Traditional elicited significantly greater workload than both Compact and Extended in all dimensions where a main effect was found: Mental Demand ($p_{\text{adj}} = .015$ and $p_{\text{adj}} < .001$), Effort (both $p_{\text{adj}} = .004$), Frustration ($p_{\text{adj}} < .001$ and $p_{\text{adj}} = .004$), and Overall Load ($p_{\text{adj}} = .007$ and $p_{\text{adj}} < .001$).

No significant differences were observed between Compact and Extended in any of these dimensions ($p_{\text{adj}} \geq .64$), suggesting comparable perceived workload between the two spatial designs.




Figure 3.5: Mean NASA TLX ratings ($\pm$ SEM) by condition.

3.6.3 User Experience

3.6.3.1 User Experience (UEQ-S)

Results for Pragmatic Quality, Hedonic Quality, and the Overall user experience score are shown in Figure 3.6. Ratings were collected on a 7-point semantic differential scale and transformed to a scale from -3 to $+3$. Values between ± 0.8 represent a mostly neutral evaluation, while higher and lower values reflect positive and negative experiences, respectively; scores beyond ± 2 are considered exceptionally rare.

Analysis revealed a significant main effect of condition on each subscale and the Overall UX score. Pairwise comparisons showed that the Traditional condition was rated significantly worse than Compact or Extended (Pragmatic, Hedonic, and Overall: all $p_{\text{adj}} < .001$). In contrast, no significant differences were found between the Compact and Extended interfaces on any dimension (all $p_{\text{adj}} > .40$).




Figure 3.6: UEQ-S results: Mean ratings ($\pm$ SEM) by condition

3.6.3.2 Evaluative User Ratings

Results and statistical comparisons for the user evaluations are summarized in Figure 3.7. These results indicate that participants rated the Traditional interface more negatively across nearly all dimensions. Significant main effects were observed in six of the seven items ($p < .05$), with *Transparency* (no main effect: $p = .079$) being the only exception. In the remaining areas, both Compact and Extended conditions generally outperformed Traditional, with significant differences found for *Distraction*, *Reorientation*, *Support*, *Disruption*, *Interruption*, and *Concentration*.

Pairwise comparisons revealed that for *Support*, *Disruption*, *Interruption*, and *Concentration*, both spatial conditions significantly outperformed Traditional after Holm corrections (all $p_{\text{adj}} < .05$). Although a significant main effect was found in *Distraction*, no individual pairwise differences were found. The only dimension showing a significant difference between the spatial interfaces was *Reorientation*, where Compact was rated significantly higher than Extended ($p_{\text{adj}} = .027$) for supporting spatial memory when returning to the document from task focus, while Extended was not rated significantly differently from Traditional ($p_{\text{adj}} = .203$). No other dimen-

sion showed a significant difference between spatial interfaces, although differences between Traditional and Extended were generally larger than those between Traditional and Compact.

Overall, the findings suggest that spatially expansive navigation supported a more focused and less disruptive experience, helping participants manage attention and reduce interruptions more effectively than with the baseline Traditional interface.




Figure 3.7: Evaluative ratings across conditions

3.6.3.3 Interface Rankings

Participants ranked all conditions from first to third across the metrics of *Speed*, *Efficiency*, *Satisfaction*, desirability for *Future Use*, and overall *Preference*. A visual representation of rankings and analysis summary is presented in Figure 3.8. Significant main effects were found for all five metrics ($p < .001$ in all cases), and post hoc tests revealed that for every metric, Traditional was rated significantly lower than both spatial conditions (all $p_{\text{adj}} < .01$). The Traditional layout was ranked lowest in all categories, receiving only a single first-place vote in three out of five metrics—all from the same participant.

These rankings reinforce findings of a preference for spatial navigation approaches over the baseline. The Compact and Extended conditions were similarly well-received, with no significant differences between them across any metric; though the Extended layout performed slightly better in *Future Use* and *Preference*, whereas Compact was favoured in *Speed* and *Satisfaction* rankings, hinting further at the

inherent trade-offs. This can also be observed in the distribution of rankings. For example, although the Extended condition received twice as many first-place rankings for *Efficiency* compared to Compact (14 vs. 7), this difference did not reach statistical significance ($p_{\text{adj}} = .137$), suggesting inconsistency in participant preference.




Figure 3.8: Post-experiment ranking by condition.

3.6.4 Summary of Findings

1. Spatially organized interfaces (*Extended* and *Compact*) yielded significantly better **performance** than the Traditional navigation methods, as measured by task completion time, total interaction count, and interaction rate.
2. Participants reported significantly lower perceived **task load** when using the spatial interfaces.
3. Spatial interfaces also resulted in significantly more positive **user experience** ratings.
4. While *Extended* and *Compact* performed **similarly** in many areas, key differences emerged:
 - (a) *Extended* required fewer **interactions** than *Compact*.
 - (b) Participants rated *Compact* as significantly more helpful for supporting **re-orientation** within the document.

3.7 Discussion and Future Work

This work offers new insights into the viability of spatial interfaces in mixed reality (MR), expands the design space of Space Filling Thumbnails (SFT) in this context, and introduces a novel study methodology for evaluating document navigation. By embedding navigation within a concurrent, hands-on task, we provide a more ecologically valid alternative to prior studies that isolate navigation as a standalone activity. Our results also support a shift away from desktop-inspired paradigms, advocating for MR-native spatial design principles. We structure the following discussion around core findings, limitations, and directions for future work.

3.7.1 Improved performance with spatial layouts

In our study, we found that completing the assembly task while navigating reference documentation with spatial layouts required significantly less time and fewer UI interactions than the traditional interface. Further, participants significantly preferred spatial layouts and rated them as more usable on a number of dimensions.

These differences are important because, unlike previous work, our measurement of performance was much coarser. Typically, previous work has required participants to repeatedly revisit parts of the document, allowing focused interactions on using spatial navigation. Our study asked participants to complete a LEGO model (for each condition), which required 9-10 steps that were described on different pages. The divided and unfocused attention in working with the spatial navigation interfaces would create increased variability in our performance data. Despite this, participants consistently and significantly performed better with the spatial layouts. We believe there are likely two main reasons for this.

First, the spatial layouts required significantly fewer interactions with the interface.

In the traditional interface, participants could scroll, use pagination buttons, or enter a page number directly. Even when they knew the desired page, these methods required multiple hand-tracked interactions—clicking buttons, scrollbars, or input fields—to navigate. While not especially frustrating, hand-tracking is inherently slower and more cumbersome than familiar tools like a mouse. Moreover, the traditional interface lacked the concrete visual cues of the spatial layouts, making it easier to overshoot a page while scrolling or mistype a page number, both of which added corrective interactions. This increased interaction burden likely contributed to the longer task times.

Second—and more importantly—participants were able to use the spatially stable overview to quickly learn and reinforce the locations of document sections. They relied on visual search to identify the desired page, which they found quick and natural. In addition to the performance results, participant ratings provide evidence for this as they found the spatial interfaces had lower task loading, were better for re-orienting themselves, were less disruptive, and better allowed them to focus on the primary task.

Together, our results strongly suggest that spatial layouts perform better than traditional interfaces. While this has been shown in past experiments, our work shows this in the context of MR-HMDs, and that these performance benefits can actually improve performance on a primary physical task.

3.7.2 Was it really spatial memory?

Human spatial memory is the underlying mechanism that has been well-established as the reason for improved performance in spatial layouts, like the space-filling thumbnails that we based our designs upon. While none of our evaluation metrics explicitly measured spatial memory, prior work on spatial interfaces has been

able to directly measure the formation of spatial memory and attribute performance gains to people using memory of locations rather than visually searching [19, 116]. Because we opted for a much larger task unit in our work (i.e., we measured performance of assembling an entire LEGO model), our study design does not allow us to definitively say that people were using this mechanism. However, through our results and observations, it seems extremely likely that people were recalling the location of pages in the document to rapidly and easily find the information they were looking for. These interfaces rely on spatial consistency to function and be used effectively, so improved outcomes imply users were able to internalize and utilize spatial information to support navigation. In further support of this, the design of decoupled and disparate information sections within the document encouraged reliance on location-based memory rather than semantic recall, supporting the notion that spatial organization facilitated navigation rather than knowledge memory [160, 116].

Further support comes from qualitative responses. When asked in the end-of-study questionnaire whether they could “remember their relative position within the document when returning to the navigation system,” participants responded positively for both spatial conditions—unanimously so for Compact, with most selecting *Strongly Agree*.

Unlike prior studies in which spatial navigation was the primary task, our design treated it as a secondary support activity within a larger workflow. We aimed to determine whether the necessary spatial memory still forms effectively under such conditions. Our findings suggest that it does: participants navigated efficiently, despite the secondary nature of the task. However, whether this memory is more fragile, less precise, or more cognitively demanding to retrieve in multitasking contexts remains unclear. While we cannot attribute these effects directly to the use of spatial

memory, previous work, as well as our results and observations, strongly suggest this was the case. Exploring and clarifying this relationship remains a valuable direction for future work.

3.7.3 Compact or Extended spatial layouts?

As previously noted, prior research on 3D spatial layouts in immersive environments has shown limited performance differences between different arrangements. Liu et al. found that a semicircular layout performed comparably to a flat layout across most metrics, with the only exception being reduced travel distance in large configurations. However, users expressed a clear preference for the curved design [138]. Follow-up work reinforced this trend, emphasizing subjective preference over measurable performance benefits in spatial memory tasks [120].

In our study, the visual distinction between the Compact and Extended layouts was somewhat more pronounced than in these prior works. *Compact* (similar to the flat layout) remained fully within the user’s field of view, while *Extended* (similar to curved layouts) leveraged peripheral space to display larger thumbnails and enhance readability. Despite this distinction, performance between the two spatial conditions was largely similar across most measures.

However, participants did differentiate between the layouts on one dimension: *Reorientation*—defined as “support for memory of relative location within the book upon return to the interface.” We hypothesized that the Compact layout might better support reorientation by presenting all pages simultaneously within a single field of view, offering a complete visual overview. In contrast, Extended placed some pages outside of the central viewport, which may have disrupted participants’ ability to maintain a holistic spatial model. Participant responses supported this hypothesis, with Compact rated significantly higher on this dimension. It is worth noting that

while our reference manual was substantial (90 pages), it remains modest in scale relative to longer documents, such as textbooks. Whether the reorientation benefit of Compact would generalize to larger documents—where further thumbnail size reductions might be necessary—is an open question for future work.

Overall, while some participants preferred Compact, more appeared to favour Extended. Given that performance was comparable and each layout offers distinct experiential benefits, it may be advantageous for designers to support both modes, allowing users to toggle based on task demands or aesthetic preference.

3.7.4 Broader applications and future directions

While our study demonstrates the benefits of spatially organized interfaces for document navigation, it focused on a single-document, single-user, task-specific context. Future research should investigate how these design strategies extend to broader document-centric activities, such as information retrieval or repeated referencing. More complex scenarios—such as multi-user environments or collaborative creative contexts like tabletop role-playing—may further highlight the strengths and limitations of spatial layouts. Indeed, the concept of spatial organization has already shown value in domains such as toolbar design for desktop word processors [161] and command layouts in mobile drawing applications [139].

There remains considerable opportunity to explore this design space. Future studies might evaluate more expansive or wraparound layouts with larger thumbnails, or investigate dynamic systems that adapt based on task complexity or user preferences. Our current evaluation was limited to a single document type and task; exploring documents of different structure, size, and semantic complexity—as well as alternative primary tasks—will help determine the generalizability of our findings.

Spatial interfaces may also offer a more intuitive and embodied way to interact with information, particularly for users with physical or cognitive differences. Such systems could reduce reliance on the frequent or high-precision interactions common in conventional document tools. Future work might explore how spatial layouts support diverse user needs by enabling alternative workflows, reducing motor demands, and enhancing spatial recall.

3.7.5 Limitations

It should be noted that our study is subject to several limitations that may affect its generalizability. First, regarding participants, our pool was relatively small, were relatively young, and more experienced with MR than the general population, given the university setting. While this aligns with prior work, future studies should target larger and more diverse samples to provide greater confidence in the findings. Second, our LEGO assembly task offered a concrete, hands-on context with divided attention, but represents only one type of workflow. Scenarios involving considerable movement, multiple users, or more complex multi-faceted activities may surface different demands and outcomes. Finally, our use of a structured, step-based manual and current MR hardware may limit transferability to denser or more visually and semantically complex documents, which may alter the balance between spatial recall, readability, and interaction burden.

3.8 Conclusion

In this work, we have demonstrated the value of spatially organized interfaces for document navigation in mixed reality. By introducing intuitive, non-linear layouts that leverage spatial memory and visual search skills, we enabled faster, more efficient

navigation with reduced task load compared to traditional approaches. Participants consistently preferred spatial interfaces, ranking them higher across all subjective measures. Notably, while the *Compact* and *Extended* layouts performed similarly across most dimensions, it seems likely that adopting both layouts would be a safe option for designers aiming to support spatial navigation. Crucially, we show that spatial memory can form and be used effectively even when navigation is a secondary task—a promising result for the real-world adoption of spatial layouts for MR workflows, such as space-filling thumbnails.

Our work supports the idea that the design of MR systems can move beyond desktop metaphors and toward interaction paradigms that embrace spatially informed design philosophies. By embedding navigation in embodied space, these systems can support more intuitive, performant, and enjoyable workflows. As the hardware and ecosystem continue to mature, the opportunity to rethink how users interact with tasks and information in 3D environments—each with their own affordances and limitations—becomes ever more compelling. This work takes a step in that direction, offering empirical evidence, practical insight, design motivation, and an invitation for further exploration, refinement, and reshaping of what mixed reality interfaces can achieve.

Chapter 4

Conclusion

4.1 Discussion & Reflection

Taken together, the two studies in this thesis offer foundational insights into the role of long-form text in spatial computing. They reveal both the opportunities and limitations of MR as a reading medium across diverse tasks and textual forms. At the same time, they raise important questions for future research: How do individual differences—such as spatial ability, visual acuity, or prior experience with MR—shape reading outcomes? What are the long-term effects of immersive reading on attention, comprehension, and fatigue? How might collaborative or social reading practices unfold in MR, and what affordances would such scenarios require?

This section offers a broader synthesis of themes that emerge when the two studies are viewed together. Although each includes its own discussion, new perspectives arise in reflecting across both—highlighting tensions between immersion and effort, novelty and familiarity, and structure and spatiality. The following subsections explore these themes and their implications for the design and understanding of reading in MR environments.

Is it the right time for these studies?

Both studies presented here faced clear limitations imposed by current-generation MR hardware—technical and ergonomic alike. Given these constraints, it is fair to ask whether it is premature to investigate reading and navigation in MR. We argue that it is not. Even if today’s devices cannot yet deliver an ideal experience, early-stage research is critical for guiding the evolution of both systems and expectations. Our first study demonstrates that reading on a current MR device may already offer distinct advantages over other digital reading options. Moreover, while MR hardware will continue to improve—becoming lighter, more comfortable, and more capable—the core challenges of head-worn displays are likely to persist. Anticipating these challenges through empirical study enables designers to shape future experiences intentionally, rather than retrofitting them post hoc.

This work provides insight into how the physical and display properties of MR systems affect reading and navigation—contexts that differ markedly from common MR use cases such as object manipulation, interactive visualization, or gaming. Each task domain demands its own design priorities. In the reading study, our findings point to the value of reflective displays and the importance of minimizing fatigue through lighter devices and better ergonomic fit. In the navigation study—while hardware comparisons were not a direct focus—we nonetheless observed issues with display fidelity, colour contrast, and device form factor. These affected the visibility of both physical and digital elements, and increased bodily movement in extended layouts. Such issues underscore the importance of accurate colour rendering, wide fields of view, and spatial coherence for reading tasks in MR.

What is a book? Or: what might a book be?

A recurring theme across this thesis is the tension between familiar reading practices and the new display and interaction possibilities enabled by MR. Much of our design effort focused on adapting conventional systems into an MR ecosystem—preserving elements like sequential order and the page metaphor—even as the medium itself affords more experimental possibilities. While this conservatism supported clarity, reduced learning demand, and comparability, it raises an important question: should MR reading interfaces emulate, or evolve?

Digitization has already blurred the boundaries of what a book is: the balance between artifact and idea, form and function. If a bound volume of blank pages can be a book, and so can a digital file, then what defines a book? Kovač et al. [162] posed this very question, critiquing legacy definitions from ISBN and UNESCO and proposing instead four defining attributes: length, textual content, formal boundaries, and information architecture. Their framework ranks “book-ness” hierarchically—placing printed novels near the top and ephemeral or boundary-blurring formats lower. In their follow-up work, *Is This a Book?* [163], Phillips and Kovač refine this into a dynamic, tripartite model—likened to a three-body problem: an open, unstable system for which no general solution exists [164]—concluding, “it is impossible to create a unique and long-lasting definition of the book” [163].

Immersive environments only heighten this instability. In theory, MR reading experiences need not adhere to traditional constraints at all: text can be layered in space, anchored to surfaces, or dynamically modulated by movement, attention, or narrative state. Physicality—whether literal or representational—can become fluid and temporal: unfolding, transforming, incorporating haptic feedback, or even vanishing entirely. These possibilities open the door to more embodied and performative modes of reading—where the user might, for instance, follow procedural steps in-place, nav-

igate fiction through movement, or uncover meaning through spatial interaction.

Yet both studies demonstrate that while readers are intrigued by novel layouts and immersive effects, certain expectations remain remarkably stable. Control, legibility, continuity, and task clarity continue to anchor user experience. In the navigation study, the Extended layout reduced necessary interactions and garnered more interest, but was still rated as less effective for recall than the more structured Compact layout. In the narrative study, immersive presentation promoted flow and engagement—until physical discomfort and visual fatigue began to undercut it. These results suggest that MR reading design must carefully balance theatricality and coherence, openness and clarity, presence and performance cost. Designing for what a book *might be* must remain attentive to what users still expect a book *to be*.

As MR books evolve beyond static emulations of physical form, designers face a critical creative space: not just preserving conventions, but deciding what to discard, reimagine, or invent. Our results highlight an intersection between narrative structure, spatial affordances, and perceptual and cognitive constraints. The future of the book in MR may not lie in any single format, but in a constellation of format-task pairings shaped by genre, context, and intent. Rather than fixate on what a book is, we might more productively ask what a book could become.

What are the demands on readers in MR environments?

Both studies share a common concern: MR reading imposes changing demands on readers' mental, physical, and attentional resources. In traditional contexts, many of these demands are absorbed by the medium itself: text appears in stable, familiar layouts, packaged into discrete physical or digital objects. In MR, however, readers must contend with additional responsibilities: managing spatial orientation,

adapting to (as yet) unfamiliar interactions, and navigating content that exists independently in space. As these studies show, transitioning reading into MR reshapes how effort is distributed—requiring readers to take on tasks once handled invisibly by the medium.

In the long-form reading study, this redistribution of effort disrupted narrative flow in subtle ways. While some burdens were reduced—such as holding a book or tuning out peripheral distractions—others emerged. Readers found themselves dividing attention between the imaginative act of immersion and the practicalities of floating text, unfamiliar input schemes, and the persistent physicality of a headset. Prior studies echo this: participants reading short passages in VR often report novelty and engagement, but also discomfort, usability issues, and attentional strain, alongside poorer comprehension scores relative to paper or screen [165]. In our study, discomfort from headsets, imprecision in hand tracking, and the lack of environmental awareness occasionally broke immersion—revealing a fundamental tension between a system’s ability to allow the user lose themselves in a book as the system itself fades into the background, and its tendency to intrude, rooted in its physicality.

In the navigation study, a parallel pattern emerged. The system eliminated some interaction demands—removing the need to flip through pages or hold the document—but introduced others. Participants had to maintain spatial awareness to keep content in view and adapt to spatially extended layouts. While MR offloaded certain tasks, it added new ones. Designs that support spatial memory (e.g., persistent anchors, predictable positioning) and reduce effortful interactions (e.g., ergonomic viewing distances, low-friction input) may help mitigate these demands—keeping attention on the content rather than the interface, and making the trade-off not only acceptable, but beneficial.

Importantly, both studies revealed domain-specific effects on task load. In the nar-

rative study, the AR headset was rated as requiring less effort than the LCD screen and showed lower overall task load than both the LCD and the VR headset—the latter scoring highest in physical demand. In the navigation study, while all conditions involved MR, those designed to expose clear spatial affordances for navigation and memory anchoring elicited lower ratings of mental demand, effort, and frustration than those that simply ported traditional interfaces. These adaptations did not eliminate task load, but they redistributed it in more effective ways.

Taken together, these findings suggest that MR shifts—not simply increases—reader workload. The demands it introduces are distinct, not necessarily greater. As MR interfaces mature and users become more practised, the friction introduced by novelty and misfit may fade. With iterative design and sustained research, MR has the potential to offload effort in new ways and support reading experiences that are immersive, efficient, and aligned with human capacities.

4.2 Conclusion

As the written word continues its long evolution through technological change, the move into Mixed Reality brings both continuity and disruption. This thesis set out to examine what reading might become in this new medium, through two complementary studies: one focused on narrative flow and immersive engagement, the other on spatial navigation in instructional documents. Together, they offer empirical insight into how MR can both challenge and support essential reading experiences—revealing trade-offs, design opportunities, and unexpected user behaviours across two distinct document types.

Our findings show that head-mounted devices can foster immersion and psychological flow in long-form reading, highlighting MR’s potential as a reading platform

despite current hardware limitations that cause fatigue and discomfort. Notably, reflective see-through AR displays offer advantages over emissive ones, combining spatial freedom and visual comfort for sustained reading. In the navigation study, spatially structured MR interfaces significantly improved performance and user experience with non-linear documents by reducing task load and enhancing navigation efficiency. Moreover, framing tasks as instructional guides demonstrated how MR layouts can support memory retention even when reading is a secondary activity.

From these findings, several design implications emerge. First, spatial stability and layout coherence are vital for both immersive and task-oriented reading. Second, interaction models must minimize friction: simplicity and transparency support both usability and absorption. Finally, while often beyond the scope of software design, hardware factors like visual fidelity, comfort, and fit remain foundational—especially for sustained or visually intensive reading. MR reading systems must treat physical and perceptual ergonomics as core design concerns.

Together, these studies contribute to a deeper understanding of how MR can reshape the form and function of books. Despite differing tasks and shared constraints, both underscore the value of hands-free interaction, spatial continuity, and human-centred adaptation. A central takeaway is the importance of *intentionality*: successful MR reading design depends not only on what is technically possible, but on careful attention to user needs and the distinct affordances of spatial media. MR reading is not simply a translation of existing forms into new hardware—but an opportunity to rethink how we engage with texts in a world unbound by flatness, fixed format, or constrained space.

Ultimately, this thesis invites us to reconsider what reading may become as it strays further from its ties to the page. As MR technologies mature, they will reshape not only how we see text but how we interact with it—and with the many purposes

books serve. In this transition, design becomes a matter of interface and also of interpretation: shaping the evolving relationship between user, text, and technology. By examining reading through the lens of spatial computing, this work lays the groundwork for more immersive, intentional, and human-centred futures of the book; whatever form it may take.

References

- [1] Paul Milgram and Fumio Kishino. A taxonomy of mixed reality visual displays. *IEICE Transactions on Information Systems*, E77-D(12):1321–1329, 1994. Manuscript received July 8, revised August 25, 1994.
- [2] Mark Billinghurst, Adrian Clark, and Gun Lee. A survey of augmented reality. *Foundations and Trends in Human–Computer Interaction*, 8(2-3):73–272, 2015.
- [3] Ronald T. Azuma. A survey of augmented reality. *Presence: Teleoperators & Virtual Environments*, 6(4):355–385, 1997.
- [4] Roger Chartier. *The Order of Books: Readers, Authors, and Libraries in Europe Between the 14th and 18th Centuries*. Stanford University Press, Stanford, CA, illustrated reprint edition, February 1994.
- [5] Robert Darnton. *The Case for Books: Past, Present, and Future*. PublicAffairs, New York, trade paperback edition, October 2009.
- [6] Marshall McLuhan, W. Terrence Gordon, Elena Lamberti, and Dominique Scheffel-Dunand. *The Gutenberg Galaxy: The Making of Typographic Man*. University of Toronto Press, Toronto, centennial edition edition, July 2011.
- [7] Thomas Francis Carter. *The Invention of Printing in China and Its Spread Westward*. Columbia University Press, New York, first edition edition, May 1925. Illustrations, plates, foldout map; based on Carter’s 1925 PhD thesis.

- [8] Elizabeth L. Eisenstein. *The Printing Press as an Agent of Change: Communications and Cultural Transformations in Early-Modern Europe*. Cambridge University Press, Cambridge and New York, illustrated paperback reprint edition, September 1980.
- [9] Ted Striphas. *The Late Age of Print: Everyday Book Culture from Consumerism to Control*. Columbia University Press, New York, first edition edition, April 2009.
- [10] N. Katherine Hayles. *How We Think: Digital Media and Contemporary Technogenesis*. University of Chicago Press, Chicago & London, illustrated paperback edition, June 2012.
- [11] Naomi S. Baron. *Words Onscreen: The Fate of Reading in a Digital World*. Oxford University Press, New York, illustrated reprint edition, February 2015.
- [12] Jenny Gabel, Melanie Ludwig, and Frank Steinicke. Immersive reading: Comparison of performance and user experience for reading long texts in virtual reality. In *Proceedings of the 2023 ACM Conference on Human Factors in Computing Systems (CHI)*, 2023.
- [13] Tilman Dingler, Siran Li, Niels van Berkel, and Vassilis Kostakos. Page-turning techniques for reading interfaces in virtual environments. In *Proceedings of the 2020 Australian Conference on Human-Computer Interaction (OzCHI)*, pages 454–461, 2020.
- [14] Marie Lebert. *A Short History of EBooks*. Project Gutenberg / NEF (University of Toronto), Toronto, ON, web edition edition, August 2009.
- [15] Terence Cavanaugh. Ebooks and accommodations: Is this the future of print accommodation? *Teaching Exceptional Children*, 35(2):56–61, November 2002.

- [16] John R. Craft. Portable document format (pdf): Standardizing document files for publication. *DESIDOC Journal of Library and Information Technology*, 23(1):19–24, January 2003.
- [17] Rufat Rzayev, Polina Ugnivenko, Sarah Graf, and Niels Henze. Reading in vr: The effect of text presentation type and location. In *Proceedings of the 2021 ACM Conference on Human Factors in Computing Systems (CHI)*, 2021.
- [18] Bilal Şimşek and Bekir Direkçi. The effects of augmented reality storybooks on student’s reading comprehension. *British Journal of Educational Technology*, 54(3):754–772, 2023.
- [19] Andy Cockburn, Carl Gutwin, and Jason Alexander. Faster document navigation with space-filling thumbnails. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’06, pages 1–10, New York, NY, USA, April 2006. Association for Computing Machinery.
- [20] Naomi S Baron and Anne Mangen. Doing the reading: The decline of long long-form reading in higher education. *Poetics Today*, 42(2):253–279, 2021.
- [21] Amy Mitchell, Galen Stocking, and Katerina Eva Matsa. Long-form reading shows signs of life in our mobile news world. *Pew Research Center*, 5(05), 2016.
- [22] M. Csikszentmihalyi. *Beyond Boredom and Anxiety*. The Jossey-Bass behavioral science series. Jossey-Bass Publishers, 1975.
- [23] Mihaly Csikszentmihalyi. The flow experience and its significance for human psychology. *Optimal experience: Psychological studies of flow in consciousness*, 2:15–35, 1988.

- [24] Mihaly Csikszentmihalyi. Flow: the joy of reading. *Applications of Flow in Human Development and Education: The Collected Works of Mihaly Csikszentmihalyi*, pages 227–237, 2014.
- [25] Tilman Dingler, Kai Kunze, and Benjamin Outram. Vr reading uis: Assessing text parameters for reading in vr. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–6, 2018.
- [26] Mark Billingham, Hirokazu Kato, and Ivan Poupyrev. The magicbook: a transitional ar interface. *Computers & Graphics*, 25(5):745–753, 2001.
- [27] Geonsun Lee, Jennifer Healey, and Dinesh Manocha. Vrdoc: Gaze-based interactions for vr reading experience. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 787–796. IEEE, 2022.
- [28] Lee Rainie, Kathryn Zickuhr, Kristen Purcell, Mary Madden, and Joanna Brenner. The rise of e-reading. *Pew Internet & American Life Project*, 2012.
- [29] Michelle Faverio and Andrew Perrin. Three-in-ten americans now read e-books. <https://www.pewresearch.org/short-read/2022/04/21/three-in-ten-americans-now-read-e-books/>, 2022. Pew Research Center.
- [30] Franziska Kretzschmar, Dominique Pleimling, Jana Hosemann, Stephan Füssel, Ina Bornkessel-Schlesewsky, and Matthias Schlewsky. Subjective impressions do not mirror online reading effort: Concurrent eeg-eyetracking evidence from the reading of books and digital media. *PloS one*, 8(2):e56178, 2013.
- [31] Nikoletta Alexandri. *The Impact of Reading Devices on Affective and Cognitive Responses to Fiction and Non-Fiction Reading*. Ph.d. thesis, Lancaster University, 2022. Doctoral dissertation.

- [32] Annika Schwabe, Lukas Brandl, Hajo G Boomgaarden, and Günther Stocker. Experiencing literature on the e-reader: the effects of reading narrative texts on screen. *Journal of Research in Reading*, 44(2):319–338, 2021.
- [33] Ana Alejandra Fuentes Cuiñas and VP Augusto. The smell of printed books: A qualitative comparison between printed and digital formats. *American Journal of Qualitative Research*, 6(1):214–255, 2022.
- [34] Turkan Ocal, Aydin Durgunoglu, and Lauren Twite. Reading from screen vs reading from paper: Does it really matter? *Journal of College Reading and Learning*, 52(2):130–148, 2022.
- [35] Hirohito Shibata and Kengo Omura. *Why Digital Displays Cannot Replace Paper: The Cognitive Science of Media for Reading and Writing*. Springer Singapore, January 2020.
- [36] Federico Pianzola, Katalin Bálint, and Jessica Weller. Virtual reality as a tool for promoting reading via enhanced narrative absorption and empathy. *Scientific Study of Literature*, 9(2):163–194, 2019.
- [37] Adam Nikola Kunzová, Adam Štěpánek, Daniel Echeverri, Danielle K. Langlois, and Simone Kriglstein. Stories with style: Narrative virtual and imaginary spaces of reading fiction in virtual reality. In *Proceedings of the International Conference on Mobile and Ubiquitous Multimedia*, pages 36–50, 2024.
- [38] Kangsoo Kim, Austin Erickson, Alexis Lambert, Gerd Bruder, and Greg Welch. Effects of dark mode on visual fatigue and acuity in optical see-through head-mounted displays. In *Symposium on spatial user interaction*, pages 1–9, 2019.
- [39] Austin Erickson, Kangsoo Kim, Gerd Bruder, and Gregory F Welch. Effects of dark mode graphics on visual acuity and fatigue with virtual reality head-

- mounted displays. In *2020 IEEE Conference on virtual reality and 3D user interfaces (VR)*, pages 434–442. IEEE, 2020.
- [40] Mary C Potter. Rapid serial visual presentation (rsvp): A method for studying language processing. In *New methods in reading comprehension research*, pages 91–118. Routledge, 2018.
- [41] Rufat Rzayev, Polina Ugnivenko, Sarah Graf, Valentin Schwind, and Niels Henze. Reading in vr: The effect of text presentation type and location. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, pages 1–10, 2021.
- [42] Seina Kobayashi, Kei Kanari, and Mie Sato. An examination of view-settings for long texts in vr reading. In *ACM SIGGRAPH 2021 Posters*, pages 1–2. ACM, 2021.
- [43] Samiha Sultana and Md Ashraful Alam. Exploring the impact of display field-of-view and display type on text readability and visual text search on large displays in virtual reality: Impact of display field-of-view and type on text readability and search in large vr displays. In *Proceedings of the 50th Graphics Interface Conference*, pages 1–11, 2024.
- [44] Regina Kaplan-Rakowski and Alice Gruber. An experimental study on reading in high-immersion virtual reality. *British Journal of Educational Technology*, 55(2):541–559, 2024.
- [45] Rufat Rzayev, Paweł W. Woźniak, Tilman Dingler, and Niels Henze. Reading on smart glasses: The effect of text position, presentation type and walking. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI ’18, pages 1–9, New York, NY, USA, 2018. Association for Computing Machinery.

- [46] Chunxue Wei, Difeng Yu, and Tilman Dingler. Reading on 3d surfaces in virtual environments. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 721–728. IEEE, March 2020. ISSN: 2642-5254.
- [47] Pei-Luen Patrick Rau, Jian Zheng, Zhi Guo, and Jiaqi Li. Speed reading on virtual reality and augmented reality. *Computers & Education*, 125:240–245, October 2018.
- [48] Jeanne Nakamura, Mihaly Csikszentmihalyi, et al. The concept of flow. *Handbook of positive psychology*, 89:105, 2002.
- [49] Csikszentmihalyi. Play and intrinsic rewards. *Flow and the foundations of positive psychology: The collected works of Mihaly Csikszentmihalyi*, pages 135–153, 2014.
- [50] Susan A Jackson and Mihaly Csikszentmihalyi. *Flow in sports*. Human Kinetics, 1999.
- [51] Jenova Chen. Flow in games (and everything else). *Communications of the ACM*, 50(4):31–34, 2007.
- [52] Julia Schöler and Sibylle Brunner. The rewarding effect of flow experience on performance in a marathon race. *Psychology of Sport and Exercise*, 10(1):168–174, 2009.
- [53] Lobna Hassan, Henrietta Jylhä, Max Sjöblom, and Juho Hamari. Flow in vr: a study on the relationships between preconditions, experience and continued use. In *Hawaii International Conference on System Sciences*, pages 1196–1205, 2020.

- [54] Jeroen S Lemmens, Fay E Clark, Xingjia Lyu, and Alex Taylor. Time flies when you're having flow: An experiment on time perception and challenge in a vr game. *Computers in Human Behavior Reports*, 18:100664, 2025.
- [55] Denise M Bressler and Alec M Bodzin. A mixed methods assessment of students' flow experiences during a mobile augmented reality science game. *Journal of computer assisted learning*, 29(6):505–517, 2013.
- [56] Daehwan Kim and Yong Jae Ko. The impact of virtual reality (vr) technology on sport spectators' flow experience and satisfaction. *Computers in human behavior*, 93:346–356, 2019.
- [57] Mihaly Csikszentmihalyi, Sami Abuhamdeh, and Jeanne Nakamura. Flow. *Flow and the foundations of positive psychology: The collected works of Mihaly Csikszentmihalyi*, pages 227–238, 2014.
- [58] Birte A. K. Thissen, Winfried Menninghaus, and Wolff Schlotz. Measuring optimal reading experiences: The reading flow short scale. *Frontiers in Psychology*, 9, 2018.
- [59] Rick Busselle and Helena Bilandzic. Fictionality and perceived realism in experiencing stories: A model of narrative comprehension and engagement. *Communication theory*, 18(2):255–280, 2008.
- [60] Jeff McQuillan and Gisela Conde. The conditions of flow in reading: Two studies of optimal experience. *Reading Psychology: An International Quarterly*, 17(2):109–135, 1996.
- [61] Bo-Ru Yang. *E-paper Displays*. John Wiley & Sons, 2022.
- [62] Shuji Fukai, Yasuhiro Niikura, Makoto Ikenaga, Ryo Hatsumi, Takeshi Nishi, Tetsuji Ishitani, Yoshiharu Hirakata, Jun Koyama, Shunpei Yamazaki, Rai

- Sato, et al. Reflective liquid crystal display without flicker for reducing eye strain. *Journal of the Society for Information Display*, 22(12):644–651, 2014.
- [63] Michael Bernard, Bonnie Lida, Shannon Riley, Telia Hackler, and Karen Janzen. A comparison of popular online fonts: Which size and type is best. *Usability news*, 4(1):2002, 2002.
- [64] DANIEL WILL-HARRIS. Georgia & verdana–typefaces designed for the screen (finally). *Typofile Magazine*. Retrieved on June, 16:1998, 1998.
- [65] Gregor Franken, Maruša Pangerc, and Klementina Možina. Impact of typeface and colour combinations on lcd display legibility. *Emerging science journal*, 4(6):436–442, 2020.
- [66] Junxiang Chen, Pengyu Zhu, Jing Yun, Baixuan Tian, Yalan Yang, and Na Liu. Text readability of smartphone in dark mode: Effects of font type, font weight and color. In *Human Aspects of IT for the Aged Population – International Conference, ITAP 2023, Held as Part of HCI International 2023, CCIS 1328*, volume 1328 of *Communications in Computer and Information Science*, pages 23–33. Springer, 2023.
- [67] Saverio Debernardis, Michele Fiorentino, and Michele Gattullo. Text readability in head-worn displays: Color and style optimization. *IEEE Transactions on Visualization and Computer Graphics*, 20(5):878–885, 2014.
- [68] Eunhee Chang, Hyun Taek Kim, and Byounghyun Yoo. Virtual reality sickness: a review of causes and measurements. *International Journal of Human–Computer Interaction*, 36(17):1658–1682, 2020.
- [69] Jacob Grimm and Wilhelm Grimm. *The complete Grimm’s fairy tales*, volume 3. Chartwell Books, 2016.

- [70] Rudolph Flesch. A new readability yardstick. *Journal of applied psychology*, 32(3):221, 1948.
- [71] John L Sherry. Flow and media enjoyment. *Communication theory*, 14(4):328–347, 2004.
- [72] Yongjun Dan. Examining the relationships between learning interest, flow, and creativity. *School Psychology International*, 42(2):157–169, 2021.
- [73] Arlen C Moller, Brian P Meier, and Robert D Wall. 9 developing an experimental induction of flow: Effortless action. *Effortless attention: A new perspective in the cognitive science of attention and action*, page 191, 2010.
- [74] Theodoros A. Kyriazos, Anastassios Stalikas, Konstantina Prassa, Michael Galanakis, Katerina Flora, and Varvara Chatzilia. The flow short scale (fss) dimensionality and what mimic shows on heterogeneity and invariance. *Psychology*, 9(6):1357–1382, June 2018. Number: 6 Publisher: Scientific Research Publishing.
- [75] Joshua Gold and Joseph Ciorciari. A review on the role of the neuroscience of flow states in the modern world. *Behavioral Sciences*, 10(9):137, 2020.
- [76] Sandra G Hart and Lowell E Staveland. Development of nasa-tlx (task load index): Results of empirical and theoretical research. In *Advances in psychology*, volume 52, pages 139–183. Elsevier, 1988.
- [77] James C Byers, AC Bittner, and Susan G Hill. Traditional and raw task load index (tlx) correlations: Are paired comparisons necessary. *Advances in industrial ergonomics and safety*, 1:481–485, 1989.

- [78] Sandra G. Hart. Nasa-task load index (nasa-tlx); 20 years later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 50(9):904–908, 2006.
- [79] Keith C Hendy, Kevin M Hamilton, and Lois N Landry. Measuring subjective workload: when is one scale better than many? *Human Factors*, 35(4):579–601, 1993.
- [80] Thomas E Nygren. Psychometric properties of subjective workload measurement techniques: Implications for their use in the assessment of perceived mental workload. *Human factors*, 33(1):17–33, 1991.
- [81] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2024.
- [82] Cassandra Phoenix and Noreen Orr. Pleasure: A forgotten dimension of physical activity in older age. *Social science & medicine*, 115:94–102, 2014.
- [83] James J Cummings and Jeremy N Bailenson. How immersive is enough? a meta-analysis of the effect of immersive technology on user presence. *Media psychology*, 19(2):272–309, 2016.
- [84] C Kuhne, ED Kecelioglu, S Maltby, RJ Hood, B Knott, E Ditton, FR Walker, and MG Kluge. Direct comparison of virtual reality and 2d delivery on sense of presence, emotional and physiological outcome measures. *frontiers in virtual reality*, 4, 2023.
- [85] Kodai Ito, Mitsunori Tada, Hiroyasu Ujike, and Keiichiro Hyodo. Effects of the weight and balance of head-mounted displays on physical load. *Applied Sciences*, 11(15):6802, 2021.

- [86] Takanori Chihara and Akihiko Seo. Evaluation of physical workload affected by mass and center of mass of head-mounted display. *Applied ergonomics*, 68:204–212, 2018.
- [87] Anna Pereira, Tevis Miller, Yi-Min Huang, Dan Odell, and David Rempel. Holding a tablet computer with one hand: effect of tablet design features on biomechanics and subjective usability among users with small hands. *Ergonomics*, 56(9):1363–1375, 2013.
- [88] Audrey Girouard, Orit Shaer, Erin T Solovey, G Michael Poor, and Robert JK Jacob. The reality of reality-based interaction: Understanding the impact of a framework as a research tool. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 26(5):1–35, 2019.
- [89] Shirley Gregor. The nature of theory in information systems. *MIS quarterly*, pages 611–642, 2006.
- [90] Katherine J Hill and Peter A Howarth. Habituation to the side effects of immersion in a virtual environment. *Displays*, 21(1):25–30, 2000.
- [91] Taylor A Doty, Jonathan W Kelly, Stephen B Gilbert, and Michael C Dorneich. Cybersickness abatement from repeated exposure to vr with reduced discomfort. *IEEE Transactions on Visualization and Computer Graphics*, 2024.
- [92] U.S. Bureau of Labor Statistics. American time use survey — 2023 results. <https://www.bls.gov/news.release/atus.nr0.htm>, June 2024. Press release.
- [93] So-Youn Jang, Jisu Park, Maria Engberg, Blair MacIntyre, and Jay D Bolter. Realitymedia: immersive technology and narrative space. *Frontiers in Virtual Reality*, 4:1155700, 2023.

- [94] Hendrik Jan Sijrier, Andrew Donlan, Shriram Pathak, Rudolf Pomme, and Erik Stephens. Importance of instruction manuals for use with hazardous area equipment. In *2009 Conference Record PCIC Europe*, pages 118–131. IEEE, 2009.
- [95] Mary Hegarty and Marcel-Adam Just. Constructing mental models of machines from text and diagrams. *Journal of memory and language*, 32(6):717–742, 1993.
- [96] James Mann Judisch, Bruce A. Rupp, and Robert A. Dassinger. Effects of manual style on performance in education and machine maintenance. *IBM Systems Journal*, 20(2):172–183, 1981.
- [97] Leon Eversberg and Jens Lambrecht. Evaluating digital work instructions with augmented reality versus paper-based documents for manual, object-specific repair tasks in a case study with experienced workers. *The International Journal of Advanced Manufacturing Technology*, 127(3):1859–1871, 2023.
- [98] Akos Nagy, Yannis Spyridis, Gregory J Mills, and Vasileios Argyriou. User experience evaluation of ar assisted industrial maintenance and support applications. In *2024 10th International Conference on Virtual Reality (ICVR)*, pages 270–277. IEEE, 2024.
- [99] Ronald T Azuma. A survey of augmented reality. *Presence: teleoperators & virtual environments*, 6(4):355–385, 1997.
- [100] Steven Henderson and Steven Feiner. Exploring the benefits of augmented reality documentation for maintenance and repair. *IEEE transactions on visualization and computer graphics*, 17(10):1355–1368, 2010.
- [101] Arthur Tang, Charles Owen, Frank Biocca, and Weimin Mou. Experimental evaluation of augmented reality in object assembly task. In *Proceedings. In-*

- ternational symposium on mixed and augmented reality*, pages 265–266. IEEE, 2002.
- [102] Book Sadprasid, Carl Gutwin, and Scott Bateman. Improving video navigation for spatial task tutorials by spatially segmenting and situating how-to videos. In *Proceedings of the 2024 Symposium on Spatial User Interaction (SUI '24)*, SUI '24, New York, NY, USA, 2024. Association for Computing Machinery.
 - [103] Oyindolapo Olabisi Komolafe, Colin A Brett, Pouya Pournasir, Lloyd Waugh, Zhen Lei, Daniel J Rea, and Scott Bateman. The effect of augmented reality on performance, task loading, and situational awareness in construction inspection tasks. In *Proceedings of the 2024 ACM Symposium on Spatial User Interaction*, pages 1–12, 2024.
 - [104] Norliza Katuk and Nur Haryani Zakaria. Linear and non-linear navigations of learning content. *International Journal of Computer Information Systems and Industrial Management Applications*, 7:10–10, 2015.
 - [105] Jesus Tejada. On paper or hypermedia? the effects of procedural information in digital video format in the learning of a music score editor program by novice users. *Revista de la Lista Electrónica Europea de Música en la Educación*, December 2007.
 - [106] Karen Schriver, Clare-Marie Karat, et al. A scoping study to explore best practice in product manual design. Technical report, UK Office for Product Safety and Standards, 2020. Accessed June 16, 2025.
 - [107] Donald Byrd. A scrollbar-based visualization for document navigation. In *Proceedings of the fourth ACM conference on Digital libraries*, pages 122–129, 1999.

- [108] Edward W Ishak and Steven K Feiner. Content-aware scrolling. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*, pages 155–158, 2006.
- [109] Carl Gutwin, Andy Cockburn, and Nickolas Gough. A field experiment of spatially-stable overviews for document navigation. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, pages 5905–5916, Denver, Colorado, USA, May 2017. ACM.
- [110] Scott Bateman, Carl Gutwin, Manasi Shah, and Nathaniel Brewer. Supporting visual search and spatial memory in a mobile application launcher. Technical Report TN18-240, Faculty of Computer Science, University of New Brunswick, 2018.
- [111] C.H. Roberts and T.C. Skeat. *The Birth of the Codex*. OUP/British Academy, 1987.
- [112] Hirohito Shibata, Kentaro Takano, and Kengo Omura. Comparison of paper and computer displays in reading including frequent movement between pages. In *Proceedings of the 26th Australian Computer-Human Interaction Conference on Designing Futures: the Future of Design*, pages 549–558, Sydney New South Wales Australia, December 2014. ACM.
- [113] Andy Cockburn, Joshua Savage, and Andrew Wallace. Tuning and testing scrolling interfaces that automatically zoom. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 71–80, 2005.
- [114] Tom Owen, George Buchanan, Parisa Eslamochilar, and Fernando Loizides. Supporting early document navigation with semantic zooming. In *The Role of Digital Libraries in a Time of Global Change: 12th International Conference*

- on Asia-Pacific Digital Libraries, ICADL 2010, Gold Coast, Australia, June 21-25, 2010. Proceedings 12*, pages 168–178. Springer, 2010.
- [115] Kenton O’Hara, Abigail Sellen, and Richard Bentley. Supporting memory for spatial location while reading from small displays. In *CHI’99 extended abstracts on Human factors in computing systems*, pages 220–221, 1999.
 - [116] George Robertson, Mary Czerwinski, Kevin Larson, Daniel C Robbins, David Thiel, and Maarten Van Dantzich. Data mountain: using spatial memory for document management. In *Proceedings of the 11th annual ACM symposium on User interface software and technology*, pages 153–162, 1998.
 - [117] Andy Cockburn. Revisiting 2d vs 3d implications on spatial memory. In *Proceedings of the Fifth Conference on Australasian User Interface - Volume 28, AUIC ’04*, pages 25–31, AUS, 2004. Australian Computer Society, Inc.
 - [118] Andy Cockburn and Bruce McKenzie. 3d or not 3d? evaluating the effect of the third dimension in a document management system. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 434–441, 2001.
 - [119] Joey Scarr, Andy Cockburn, Carl Gutwin, et al. Supporting and exploiting spatial memory in user interfaces. *Foundations and Trends® in Human-Computer Interaction*, 6(1):1–84, 2013.
 - [120] Jiazhou Liu, Arnaud Prouzeau, Barrett Ens, and Tim Dwyer. Effects of display layout on spatial memory for immersive environments. *Proceedings of the ACM on Human-Computer Interaction*, 6:468–488, 2022.
 - [121] Johannes Zagermann, Ulrike Pfeil, Daniel Fink, Philipp von Bauer, and Harald Reiterer. Memory in motion: The influence of gesture-and touch-based input

- modalities on spatial memory. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 1899–1910, 2017.
- [122] Tallie Z Baram, Flavio Donato, and Gregory L Holmes. Construction and disruption of spatial memory networks during development. *Learning & memory*, 26(7):206–218, 2019.
- [123] B. Kolarik and A. Ekstrom. *The neural underpinnings of spatial memory and navigation*, volume 3, pages 507–514. Elsevier Inc., 2015.
- [124] Silvia Serino, Daniele Di Lernia, Giulia Magni, Paolo Manenti, Stefano De Gasperi, Giuseppe Riva, and Claudia Repetto. Egocentric and allocentric spatial memory for body parts: a virtual reality study. *Journal of Cognition*, 7(1):33, 2024.
- [125] Md Sami Uddin and Carl Gutwin. The image of the interface: How people use landmarks to develop spatial memory of commands in graphical interfaces. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–17, 2021.
- [126] Bryony Olney and Bobby Nisha. Spatial cognition with virtual reality. *Conscious Cities: Bridging Neuroscience, Architecture, and Technology*, 2:142–147, May 2017.
- [127] Neil Burgess. Spatial memory: how egocentric and allocentric combine. *Trends in cognitive sciences*, 10(12):551–557, 2006.
- [128] Masood Fathi, Ingemar Karlsson, Göran Grahñ, and Andreas Björnsson. Unveiling the potential of mixed reality: enhancing time measurement and operator support in manual assembly processes. *Procedia Computer Science*, 232:2670–2679, 2024.

- [129] Antonio Maffei, Michela Dalle Mura, Fabio Marco Monetti, and Eleonora Boffa. Dynamic mixed reality assembly guidance using optical recognition methods. *Applied Sciences*, 13(3):1760, 2023.
- [130] Lei Hou, Xiangyu Wang, and Leonhard Bernold. Using animated augmented reality to cognitively guide assembly. *Journal of Computing in Civil Engineering*, September 2013.
- [131] Bernd Schwald and Blandine De Laval. An augmented reality system for training and assistance to maintenance in the industrial context. *Journal of WSCG*, 11(1–3):425–432, 2003.
- [132] Balasaravanan Thoravi Kumaravel, Fraser Anderson, George Fitzmaurice, Bjorn Hartmann, and Tovi Grossman. Loki: Facilitating remote instruction of physical tasks using bi-directional mixed-reality telepresence. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, pages 161–174, 2019.
- [133] Lili Wang, Xiangyu Li, Jian Wu, Dong Zhou, Im Sio Kei, and Voicu Popescu. Avicol: Adaptive visual instruction for remote collaboration using mixed reality. *International Journal of Human–Computer Interaction*, 41(2):1260–1279, 2025.
- [134] Zhen Li, Michelle Annett, Ken Hinckley, Karan Singh, and Daniel Wigdor. Holodoc: Enabling mixed reality workspaces that harness physical and digital content. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI ’19, pages 1–14, New York, NY, USA, May 2019. Association for Computing Machinery.
- [135] Chen Chen, Cuong Nguyen, Jane Hoffswell, Jennifer Healey, Trung Bui, and Nadir Weibel. Papertoplace: Transforming instruction documents into spatial-

- ized and context-aware mixed reality experiences. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–21, 2023.
- [136] Barrett Ens, Juan David Hincapié-Ramos, and Pourang Irani. Ethereal planes: a design framework for 2d information space in 3d mixed reality environments. In *Proceedings of the 2nd ACM symposium on Spatial user interaction*, pages 2–12, Honolulu Hawaii USA, October 2014. ACM.
- [137] Nooree Na, Kyeong Ah Jeong, and Hyeon-Jeong Suk. Do curved displays make for a more pleasant experience? In *Human vision and electronic imaging XX*, volume 9394, pages 421–426. SPIE, 2015.
- [138] Jiazhou Liu, Arnaud Prouzeau, Barrett Ens, and Tim Dwyer. Design and evaluation of interactive small multiples data visualisation in immersive spaces. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 588–597, 2020.
- [139] Carl Gutwin, Andy Cockburn, Joey Scarr, Sylvain Malacria, and Scott C Olson. Faster command selection on tablets with fasttap. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2617–2626, 2014.
- [140] Joey Scarr, Andy Cockburn, Carl Gutwin, and Andrea Bunt. Improving command selection with commandmaps. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 257–266, 2012.
- [141] Julian Rasch, Matthias Wilhalm, Florian Müller, and Francesco Chiossi. Ar you on track? investigating effects of augmented reality anchoring on dual-task performance while walking. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–21, 2025.

- [142] Donghao Ren, Tibor Goldschwendt, YunSuk Chang, and Tobias Höllerer. Evaluating wide-field-of-view augmented reality with mixed reality simulation. In *2016 IEEE Virtual Reality (VR)*, pages 93–102. IEEE, 2016.
- [143] William MacNamara. Evaluating the effectiveness of the gestalt principles of perceptual observation for virtual reality user interface design. M.sc. thesis, Technological University Dublin, 2017. Available via TU Dublin digital repository.
- [144] Niklas Stein, Tamara Watson, Markus Lappe, Maren Westendorf, and Szonya Durant. Eye and head movements in visual search in the extended field of view. *Scientific Reports*, 14(1):8907, 2024.
- [145] Shaimaa Masarwa, Olga Kreichman, and Sharon Gilaie-Dotan. Larger images are better remembered during naturalistic encoding. *Proceedings of the National Academy of Sciences*, 119(4):e2119614119, 2022.
- [146] Edgar Chan, Oliver Baumann, Mark A Bellgrove, and Jason B Mattingley. From objects to landmarks: the function of visual location information in spatial navigation. *Frontiers in psychology*, 3:304, 2012.
- [147] Norman G Vinson. Design guidelines for landmarks to support navigation in virtual environments. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 278–285, 1999.
- [148] BoYu Gao, Zhiguo Chen, Xian Chen, Huawei Tu, and Feiran Huang. The effects of audiovisual landmarks on spatial learning and recalling for image browsing interface in virtual environments. *Journal of Systems Architecture*, 117(C):102096, 2021.
- [149] Nicholas F Polys, Seonho Kim, and Doug A Bowman. Effects of information layout, screen size, and field of view on user performance in information-rich

- virtual environments. In *Proceedings of the ACM symposium on Virtual reality software and technology*, pages 46–55, 2005.
- [150] Jing Feng and Ian Spence. Attending to large dynamic displays. In *CHI'08 Extended Abstracts on Human Factors in Computing Systems*, pages 2745–2750. 2008.
 - [151] Kevin Arthur. Effects of field of view on task performance with head-mounted displays. In *Conference companion on human factors in computing systems*, pages 29–30, 1996.
 - [152] JJ-W Lin, Henry Been-Lirn Duh, Donald E Parker, Habib Abi-Rached, and Thomas A Furness. Effects of field of view on presence, enjoyment, memory, and simulator sickness in a virtual environment. In *Proceedings ieee virtual reality 2002*, pages 164–171. IEEE, 2002.
 - [153] Amit P Sawant and Christopher G Healey. A survey of display device properties and visual acuity for visualization. *North Carolina State University. Department of Computer Science. Raleigh, NC*, pages 27695–8207, 2005.
 - [154] Andy Cockburn, Amy Karlson, and Benjamin B. Bederson. A review of overview+detail, zooming, and focus+context interfaces. *ACM Computing Surveys*, 41(1):1–31, 2009.
 - [155] Stephen Fitchett, Andy Cockburn, and Carl Gutwin. Improving navigation-based file retrieval. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2329–2338, 2013.
 - [156] BrickLink. Studio 2.0. <https://www.bricklink.com/v3/studio/download.page>, 2025. Version 2.0, accessed May 2025.

- [157] Martin Schrepp, Andreas Hinderks, and Jörg Thomaschewski. Design and evaluation of a short version of the user experience questionnaire (ueq-s). *International Journal of Interactive Multimedia and Artificial Intelligence*, 4(6):103–108, 2017.
- [158] Maria Rauschenberger, Martin Schrepp, Manuel Cota, Siegfried Olschner, and Jörg Thomaschewski. Efficient measurement of the user experience of interactive products. how to use the user experience questionnaire (ueq). example: Spanish language version. *International Journal of Interactive Multimedia and Artificial Intelligence*, 2:39–45, March 2013.
- [159] Daniel Gaspar-Figueiredo, Silvia Abrahão, Emilio Insfrán, and Jean Vanderdonckt. Measuring user experience of adaptive user interfaces using eeg: a replication study. In *Proceedings of the 27th International Conference on Evaluation and Assessment in Software Engineering*, pages 52–61, 2023.
- [160] Don Norman. *The design of everyday things: Revised and expanded edition*. Basic books, 2013.
- [161] Andy Cockburn, Carl Gutwin, and Saul Greenberg. A predictive model of menu performance. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 627–636, 2007.
- [162] Miha Kovač, Angus Phillips, Adriaan van der Weel, and Rüdiger Wischenbart. What is a book? *Publishing Research Quarterly*, 35(3):313–326, September 2019.
- [163] Angus Phillips and Miha Kovač. *Is This a Book?* Elements in Publishing and Book Culture. Cambridge University Press, Cambridge, UK & New York, NY, illustrated paperback edition, June 2022.

- [164] Mauri J. Valtonen and Hannu Karttunen. *The Three-Body Problem*. Cambridge University Press, Cambridge, UK, illustrated edition edition, March 2006.
- [165] Murat Çoban, Arif Cem Topuz, Ahmet Karabulut, and Thomas K. F. Chiu. Comparing realities: A study on the impact of virtual reality versus paper-based reading on higher education students' comprehension skills. *Participatory Educational Research*, 11(4):19–36, July 2024.

Vita

Candidate's full name:

Colin Andrew Brett

University attended (with dates and degrees obtained):

2016-2021: Bachelor of Arts (Computer Science)

Memorial University of Newfoundland

Publications:

- Komolafe, O. O., Brett, C. A., Pournasir, P., Waugh, L., Lei, Z., Rea, D. J., & Bateman, S. (2024, October). The Effect of Augmented Reality on Performance, Task Loading, and Situational Awareness in Construction Inspection Tasks. In *Proceedings of the 2024 ACM Symposium on Spatial User Interaction* (pp. 1-12).
- Brett, C. A., & Bateman, S. (2025, November). Spatially Organized Interfaces for Mixed Reality Document Navigation. In *Proceedings of the 2025 ACM Symposium on Spatial User Interaction* (accepted, in press).
- Brett, C. A., & Bateman, S. (2025). Mixed Reality and the Long-Form Reading Experience [Poster]. To be presented at *the 2025 ACM Symposium on Spatial User Interaction*, Montreal, Canada, November 10-11, 2025.