Academic Reading: Technologies and Practices

by

Velian Pandeliev

A literature review submitted
in partial fulfillment of the requirements
for the degree of Doctor of Philosophy
Graduate Department of Computer Science
University of Toronto

Abstract

Academic Reading: Technologies and Practices

Velian Pandeliev
Doctor of Philosophy
Graduate Department of Computer Science
University of Toronto
2013

Reading in academic contexts usually involves a form of responsive or active reading that relies on companion activities such as annotation, note-taking, and consulting multiple documents simultaneously. Given the increase in availability, affordability, and capabilities of digital devices, multiple attempts are being made to study and improve digital academic reading, aiming for the academic reading equivalent of the paperless office. However, most of these attempts have focused on designing or evaluating single devices that are envisioned as replacements both of paper and of readers' other digital devices in the realm of academic reading. We outline the two significant obstacles that have limited the success of this impositional approach. Firstly, it is nearly impossible to create a single device that satisfies the full range of requirements for reading and its companion activities. Secondly, academic readers already employ successful reading strategies that integrate the digital devices they own into their workflow according to their strengths, and this emergent behaviour is highly individualized and incompatible with an imposed solution. We hypothesize that individual readers' experiences with digital technologies can yield valuable insight into the process of academic reading that is not available through research on technology imposition. We argue that academic reading is a multi-document, multi-technology activity that seamlessly interleaves reading (content consumption) with its companion activities (content creation), while utilizing the strengths of paper and various digital devices. Our research on academic reading will focus on understanding emergent reader strategies and facilitating integration of heterogeneous devices and technologies into a coherent workflow. As a first step, we are conducting a study to gain insight into academic readers' technology use and habits to inform our research direction.

Contents

Chapter 1

The Nature and Cognition of Reading

To many scholars, language is the vehicle not only for human communication, but also for human reasoning and cognition. Despite conflicting theories on the origins of spoken language, we know that humans have been communicating verbally for at least 100 000 years. In contrast, the first written system capable of representing language was invented in Sumer around 3200 BCE.

It may be difficult for modern humans to appreciate the technological leap that written language represents. For over 90% of the history of the thinking human, every thought, invention or communication persisted only as long as the memory of the last person who heard it described. Knowledge was highly localized and passed orally from generation to generation. The invention of writing meant that human ideas could be preserved in materials more durable than the human body, they could resist the decay of retelling or recall, and they could be multiplied, disseminated and organized. Arguments could be longer and more complex than a listener's aural attention span could support, and the thoughts and ideas of writing humans could persist for millennia. As a cognitive artifact, written text unburdens our minds from a tremendous amount of memorization and refactors the task of understanding language into a visuo-tactile one.

Reading is the complement activity to writing, the task of decoding written symbols into a stream of language in our minds. Once a skill reserved exclusively for the upper echelons of scholars, clergy, and noblemen, at present reading is an expected skill for most 7-year olds throughout the world. With the advent of industrialization and, recently, the shift towards a knowledge-based economy, worldwide literacy policies have made reading a crucial part of education. In the 20th century, scholars have been able to classify and differentiate between different types of reading based on the skill level required of the reader, the purpose of the reading task and the strategies employed.

Knowledge workers and academics engage in a process known as responsive reading, which is one of the most complex and difficult types of reading, involving deliberate training, meta-cognition and a host of additional skills. While digital technology is matching the abilities of paper books for leisure reading, the set of requirements for responsive reading is more extensive.

1.1 The cognitive underpinnings of reading

Reading is a unique activity. It requires the use of all the cognitive faculties involved in language, which has long been considered the benchmark for human intelligence, combined with visual processing, memory retrieval and inference generation. This makes it one of the most complex cognitive activities that humans engage in. It is also the primary means by which we acquire complex knowledge, and is "essential for successful participation in our society" [81]. As such, much research has been conducted on the mechanisms by which reading occurs in the mind and on our ability to comprehend written materials. The briefest of overviews of this research follows, with a discussion of memory, reading comprehension, and coherence.

1.1.1 Memory

The essentials of human memory factor greatly into our understanding of reading and its related processes. According to cognitive scientists, there are three major types of memory that differ in capacity, level of processing and permanence.

Sensory Memory

The least cognitively demanding type of memory is called *sensory memory* (SM). Sensory memory is the buffer for perception. Its three main components are iconic (visual), echoic (aural) and haptic memory. Recall durations range from 100 ms for iconic to 1000-5000 ms for echoic memory. Sensory memory applies no processing or reasoning to the information it stores and it must be consciously transferred to working memory in order to be retained or reasoned about.

For clarity, this document will not employ the term *short-term memory* as it is highly overloaded and ambiguous. Different sources use the term to refer to sensory memory, working memory, or even aspects of long-term memory at different times, rendering it essentially useless.

Working Memory

Working memory (WM) is the intermediate memory store between the extremely short-lived sensory memory and long-term memory, which can last a lifetime and has a theoretically unlimited capacity. It serves as an intermediary, enabling the simultaneous activation of several sensory or long-term concepts or images. The term "working" hints that this structure does more than just store data. Concepts stored in working memory can be combined, manipulated and altered before they are stored in long-term memory, which means that working memory is also the staging area for reasoning.

According to Baddeley [6], WM consists of three distinct specialized structures controlled by a central executive. These structures are the auditory loop, the visuospatial sketchpad and the episodic buffer. The visuospatial sketchpad is used to mentally manipulate pictures and objects. It plays a part in how readers remember the location of text on a page and in maintaining their sense of place within the text. The auditory loop, the mind's inner monologue, is the buffer we use to hum a catchy song or repeat a friend's phone number until we memorize it. It is single-channel, linear, susceptible to interruption, and low-bandwidth. When we read, we convert symbols into an aural stream of language in the mind, engaging the auditory loop to "sound out" what is being read in order to process the word being decoded and to consolidate it into the context of related terms and concepts in WM. The episodic buffer is a recently discovered structure that combines visual, auditory and temporal information to enable reasoning about causality and the sequence of events. Working memory is akin to random access memory in a computer, and it has two important properties.

The first is that working memory has a relatively low capacity. While millions of pieces of information can be stored in long-term memory, manipulating them in WM requires that they be activated or primed for reasoning. This limitation was first proposed by Miller [54] in 1956, who posited a WM capacity of approximately seven distinct chunks. Further research has found that number to vary based on the format, length, and familiarity of the stimuli [30]. Ultimately, Cowan proposed that only four or so concepts may be the immediate focus of reasoning and attention [19]. Given the fact that reading often draws on a multitude of related concepts, Ericsson and Kintsch [23] posit that the few activated concepts in working memory act as points of access or hooks for a network of related entities that reside in long-term memory, but are much easier to retrieve, creating a section of long-term memory that is primed for working memory access. Additionally, concepts modified through reading and reasoning in WM can be re-encoded into long-term memory, updating the internal representation of the text.

The second property of WM is that it requires attention not only to manipulate, but also to maintain active representations. Evidence [8] shows that mental representations in WM decay over time if they are not attended to. The central executive system that decides how to allocate attention is responsible for processing the current task, rehearsing concepts to prevent decay, and tuning out distractions. This system is serial and has a limited capacity which varies from person to person but represents a significant constraint on the amount of information that can be attended to at any moment [81]. Furthermore, that amount is lower if the task includes added complexity such as navigation, mode switching, or some form of distraction.

Long-term Memory

Long-term memory (LTM) is a structurally distinct [90] type of memory responsible for storing information that is not currently used for reasoning. If WM is akin to RAM, LTM's equivalent in a computer is the hard disk. Unlike digital storage, LTM has a potentially unlimited capacity and a duration that often matches the human lifespan. LTM can be thought of as an interconnected graph of concepts, memories and relationships that are more easily recalled if they or related concepts have been recently accessed in the past. It, too, has different components. Explicit or declarative memory stores concepts we are aware of knowing, and is divided into episodic memory (memory for events) and semantic memory (memory for abstract concepts). Implicit or procedural memory holds learned skills such as tying one's shoes, writing or riding a bicycle.

An important feature of long-term memory is that it does not hold perfect representations of stimuli or concepts. With the small exception of deliberately memorized text, what is stored in long-term memory are sparse representations of events, concepts and images, not the source stimuli themselves. Only salient details are encoded in LTM while any other information is filled in after retrieval based on common knowledge and previous experience to complete a coherent mental picture. On occasion, some of these false details that form part of the recalled memory even though they were not "loaded" from LTM can be re-encoded into LTM, leading to the real danger of changing memories through recall.

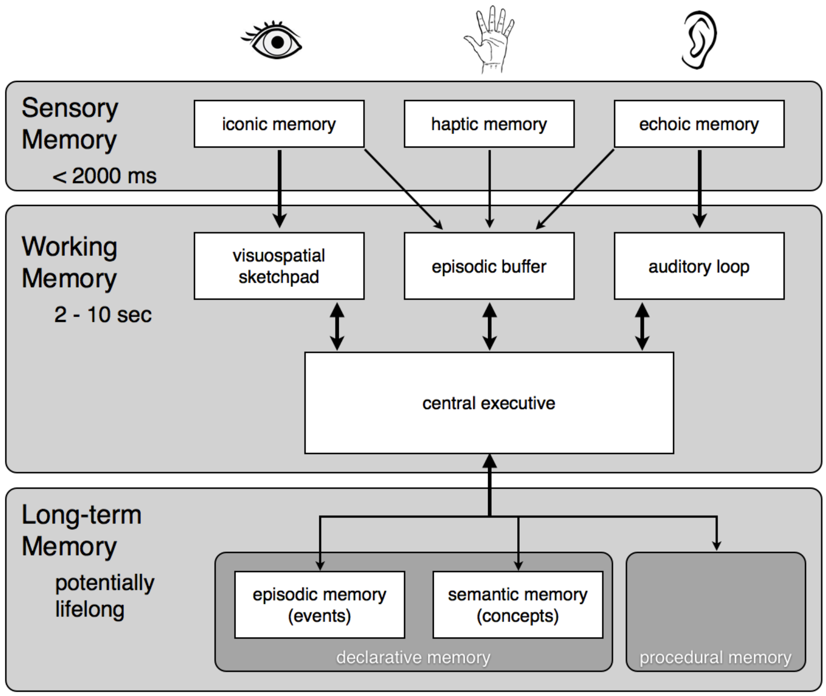

Figure 1.1: The three types of human memory.

1.1.2 Reading Comprehension

The description of memory structures above hints at a prevailing current of information flow during reading that begins at sensory memory, consolidates with recently read concepts in working memory and is learned by being encoded to long-term memory. This view was supported initially by bottom-up theories of reading [70] in which the perceptual stimulus was processed in its entirety before being sent to working memory. These theories fail to explain evidence such as the part-of-speech similarity between a mis-read word and the intended word [37] and the decrease in time required to recognize a word that fit within the reading context [80].

The current prevailing theory is an interactive model of information flow: sensory input from the reading is being combined in working memory with activated lexical, semantic, syntactic and orthographic information from LTM, in order to understand the text.

Reading comprehension refers precisely to the ability to do so: to extract the relevant information from a text and to create and remember a usable mental representation of this information for future use. It "require[s] that readers construct a "mental picture" of the text: a representation in memory of the textual information and its interpretation. Such mental representations ideally can be easily accessed, manipulated, and applied to any number of situations. Thus they are central to performance on any comprehension task. When reading is successful, the result is a coherent and usable mental representation of the text." [82]. This means that, in an ideal case, the reader has a flexible, custom-built mental data structure that enables reasoning and that can be stored in memory and used after the source text is no longer available.

It must be noted that the criteria for success, coherence, or usability here only concern the interactions between a reader, his or her mental representation of the text, and his or her goals. These descriptions assume no absolute truth or validity.

However, unlike literary prose, in which different interpretations are expected and encouraged, in technical or academic prose there is an additional criterion, which denotes how well the reader's representation matches a canonical agreement on the facts of the discipline. Given that the presumption that such an agreement exists is naive and optimistic, the criterion is perhaps more functional, concerning whether the reader's representation can successfully guide him or her through evaluations of his or her knowledge on a topic, or to generate a solution to a particular problem.

More than any other aspect of reading, comprehension is the principal, measurable criterion that reflects the central purpose of the activity.

1.1.3 Generations of Reading Comprehension Research

Research on reading comprehension has developed in three consecutive *generations*, each with a slightly different focus. As in real life, subsequent generations have not superseded the findings of the previous ones but rather opened new avenues of contribution to the field. The progression of reading comprehension research began with explorations of the nature of the successful representations of text, shifted to include the on-line nature of reading by exploring the limitations of the processing taking place during the reading activity and have expanded to define holistic computational models of reading comprehension that incorporate on-line and off-line mechanisms for concept representation, activation and reinforcement. These three generations are summarized in greater detail by van den Broeck et al. [81]

Mental Representation of Text

The *first generation* of reading comprehension research determined that the structure of the mental representation of a text is non-trivial and encodes more and simultaneously less information than is available in the original text. The encoding is a combination of a salient subset of details (the ones that the reader attended to due to attentional limitations and his or her own goals) and information not present in the text such as related concepts in background knowledge or previously known templates for its expected structure. If the reader has no prior knowledge or relevant concepts, the representation would be one of the text alone ("the text base") [27]. If the text is part of a familiar domain or closely linked to other recently read documents, its representation includes hooks to related concepts and is akin to a more general representation of a domain ("situation model") [27]. Reading goals and motivations affect where a text representation will fall along this continuum. A mental representation of a text is deemed sufficiently detailed and usable if it achieves **coherence** in two significant ways.

First, the representation must fit reasonably well within a canonical schematic structure, that is, be similar enough to an expected schema for the contents of the text. This is called **schematic coherence**. Schema theory states that a reader instantiates a template for the context of what they are reading with a set of expectations called schemata (singular schema). As they read, they engage in populating the schema, filling in the empty slots it presumes are relevant with information from the text [4]. As an example, a fictional story might have a linear chronological progression, and its schema may contain the expectation for a protagonist, a conflict, as well as the appropriate stages of the story's development including rising action, climax, etc.

According to Rumelhart [66], reading comprehension is greatly facilitated by the correct schema being activated and called to working memory so it can be completed. This is a top-down process, with concepts and interactions being slotted into a previously existing high-level structure. Schematic coherence may not be achieved if the reader has no schema for the context of the text, in which case a new one must be created with the reader having very little structure to rely on. It may also fail if a schema for the type of text exists, but it is not activated by the text, or if the text is matched to a schema that is not the one intended by the author [66].

Interestingly, a partial match to a schema leads to the reader both having some experience to draw from and being able to update his or her template, leading to improved understanding of the topic. "When the fit between old in-head information and new on-the-page information is good but not perfect, learning from the text [occurs]" [27] and schemata are updated or generated.

The domain of "technical prose" [27], i.e. prose designed to impart new information in a dense manner, poses challenges to schematic coherence. Readers may be able to use schemata for the structure of the text, recalling previous documents in a similar format, but the schemata for new content would have to be generated through inference and careful reading.

As well as achieving schematic coherence, a mental representation of text must also be relationally coherent: it must represent the relevant relationships between the elements in the text without leaving unexplained gaps or illogical leaps. These relationships are best captured as a graph of interconnected concept nodes and edges that represent the semantic relationships between them, which may be causal, emotional, spatial, referential, etc. [81] **Relational coherence** is judged by the reader and is a bottom-up process, with each new piece of information encountered in the text either answering a question about a previously unclear relationship or introducing a new relationship that must be understood before coherence can be achieved.

Research indicates that causal relationships are always encoded, while the selective inclusion of other types of relationships depends on the reader's goals, strategies, previous experience, attention and the genre of the text, which collectively form his or her *standards for coherence* [83]. These standards constitute reader-specific and situation-specific self-assessment criteria for achieving comprehension.

Limitations Of Cognitive Processes During Reading

With the advent of computerized time-sensitive research apparatus in the 1980's, more became known about the cognitive processes that took place in the course of reading. These processes are referred to as on-line as they describe the actual processes that take place while the reading is occurring, as opposed to any off-line processing or consolidation of the information from the text that might occur afterwards.

In order to generate a coherent mental representation, the reader must make inferences about the relationships between the elements in the text, and most of these inferences occur on-line, in working memory. Due to the limited capacity of WM [33] and the limits in attention allocation [31], "at any point during the reading [a reader] can only attend to a small subset of all elements that are relevant to the text. As a consequence, the reader must wisely select which elements to attend to" [81] to achieve reading comprehension.

Due to the bottleneck created by attention allocation, readers typically only generate inferences that are essential to creating or maintaining coherence in the mental representation of the text according to their own standards for coherence. These inferences often include concepts outside the text that is being read, in which case they may be supplied by the episodic memory of having read previous parts of the text or by the semantic memory for previously acquired knowledge [81], which includes previously read documents.

It is important to note that readers of the same text will form different internal representations of it since they may need the representation for vastly different purposes. Also, every representation of a text will be imperfect because it is limited by the reader's attentional resources, and it will not contain all the information available in the text.

The Landscape Model of Reading Comprehension

The third generation of reading research aims to arrive at a computational model for the mental representation of text that can integrate the results achieved by the previous two generations into a predictive theory. One of the most successful models to emerge from this generation is the landscape model of reading comprehension [84]. This model assumes a node-and-edge representation of concepts in memory akin to a neural network, with weighted connections that indicate not only the kind of relationship between concepts, but also the relevance that they have to each other. It considers the entire network as a mental landscape, and notes that at any given moment, certain concepts are activated to a greater degree than others. From time step to time step, as a text is being read, these activations change. Some concepts are activated because they are being read about, some because they are part of the reading context or domain, and some are recalled from LTM in order to fill a coherence gap for the reader. If a concept is activated, the activation propagates to its connected nodes proportionally to the strength of the connection. As with all activations in WM, these decay over time if not rehearsed or strengthened by additional activation.

When a document is being read, sensory memory holds the most recently read clause or sentence. We refer to the reading of a single clause or sentence as a cycle [22]. Given this definition, van den Broek et al. [84] identify four potential sources of concept activation:

- The current cycle of text, which includes concepts and references found in the immediate clause being read, which are moved from sensory memory into working memory for processing.
- The previous cycle of text, which has just been processed in WM and whose activations still persist.
- The rest of the text that has already been read. The activations from this part have decayed out of working memory, but the act of reading it has been encoded in episodic long-term memory and it can be recalled, reactivating relevant concepts more easily.
- Background knowledge, which is stored in semantic long-term memory. This source is consulted if the text does not hold all the necessary information to form a relationally coherent representation.

Readers rely on two different mechanisms to access the last two activations, which are not readily available in WM. One is *cohort activation*, which posits that activating one concept in memory immediately activates all the concepts adjacent to it in the graph. These activations decay rapidly, but may persist if activated again soon afterwards or focused on. The other is *coherence-based activation*, which occurs when the reader identifies a gap or lapse in his or her ability to create a coherent representation, leading to a strategic, deliberate attempt to find the information necessary to maintain coherence. To achieve this, readers may look back in the text, re-read previous sentences or chapters, or consult related material. There is a crucial difference between these two mechanisms: while cohort activation is automatic and effortless, filling a coherence gap is slow and deliberate, requiring the seeking of missing information elsewhere in the document, in memory, or in other documents. Furthermore, it necessitates a level of task monitoring and reasoning about one's performance on the task.

1.1.4 Metacognition and Executive Control

Readers are active learners that follow standards of coherence and strategically allocate resources while reading. A valuable aspect of our understanding of reading concerns how readers make decisions when reading: how they assess their own understanding, how they detect failures and what strategies they apply to compensate for them. In short, we need to consider readers' metacognition: their knowledge of their own cognitive processes and their properties and products. It is through metacognitive reasoning that a reader decides how carefully to read a passage. It is the feedback loop, the internal assessment mechanism by which we determine our own reading success.

Theorists who view the mind as an information processing system akin to a computer emphasize the importance of a related faculty: executive control. In contrast to the metacognitive view, which is concerned with the reader's knowledge, awareness and conscious access, the proponents of executive control are concerned with the symbol-manipulation aspects of the mind-computer and its ability to execute sequences of steps while allocating limited resources in the best manner possible. It is the executive control structure that enables a reader to proceed through an easy, leisurely reading quickly until an unfamiliar concept or idea is encountered, at which point additional resources may be allocated and a new approach devised.

Although borne out of conflicting philosophical views of the mind, both declarative knowledge about the process of reading and the notion of resource allocation and sequential action plans have a place in reading discourse.

Flavell [25] posits a comprehensive theory for the role of metacognitive knowledge in detecting failure while performing a task and triggering appropriate strategies to mitigate that failure.

When applied to reading, metacognitive knowledge is knowledge about one's own experience as a reader, about the task at hand and about the necessary ways in which particular documents must be read and learned. A metacognitive experience, as dubbed by Flavell, or episode, occurs when a reader applies his or her knowledge to a specific reading task. It may occur before a document is read, perhaps to signal the urgency or difficulty of the task ahead. It may occur after the reading has finished as a way of evaluating progress towards a larger goal. Most importantly, a metacognitive episode that occurs during reading is the result of the reader successfully monitoring his or her progress and detecting a cognitive failure: confusion or uncertainty about the text that prevents the reader from achieving coherence. Cognitive failures that are successfully detected and result in metacognitive reflection are considered metacognitive successes [27]. Ones that are undetected are both cognitive and metacognitive failures: the failure to understand the text has been compounded by the failure to detect the lapse and cannot be remedied. These may occur on account of poor metacognitive knowledge, or as a result of overloading the executive control in WM with superficial tasks which would diminish the availability of metacognitive resources to monitor the reading task. At the moment of a metacognitive episode, the executive control begins searching for an appropriate alternative strategy to achieve the reading goal. Reading strategies are discussed in greater detail in Section 1.2.3.

1.2 Reading types, scenarios and strategies

In this section we present a nuanced description of reading along three distinct dimensions: the exact type of activity undertaken under the umbrella term "reading", the scenario that motivated the reading and the metacognitive awareness of the strategies employed by the reader to facilitate his or her understanding.

1.2.1 Types of reading

The following distinct activities may be described as different reading types [76]. Although all involve decoding symbols, they differ greatly in the procedural sequence of steps a reader would undertake to complete the reading task.

- **Skimming** is sequential and rapid-access, often relying on section headings and illustrations to get an initial idea of the text, to determine its relevance and to decide whether a more thorough reading is necessary.

- **Scanning** involves locating occurrences of a particular search phrase which is explicitly known to the reader. This is a low-level, pattern-matching activity.

- **Search reading** is similar to scanning, but the exact written form of the concept sought is unknown, leading to reading at a higher level and matching against concepts rather than word patterns.

- **Receptive reading** is sequential, slow, and involves reading a particular passage from beginning to end as a coherent whole. It closely approximates listening behaviour [58] and is the reading activity most naively associated with the practice of reading, the one that typically occurs with fiction and pleasure reading, and the one commercial e-readers are best suited for.

- **Responsive reading** [76] is a reading type typically engaged in by knowledge workers, whereby information from the text needs to be incorporated into an existing body of knowledge in a particular domain. As the most complex reading type, it may involve episodes of the other types of reading. It is detailed in section 1.4.

In addition to the different types of reading, reading tasks often include information processing activities that cannot be classified as reading activities but are nevertheless essential to the reading process. These companion activities such as navigation, orientation, annotation, summarization and cognitive mapping [76] are discussed in greater detail in Chapter 3.

1.2.2 Reading Scenarios

Lorch et al. [40] distinguish between different kinds of reading based on the situation, the goals and the reasons why a particular book, paper or passage is being read. They create a topology of reading scenarios with respect to college students, broadly separated into reading for school, which is often assigned and evaluated in some manner and reading by choice, which is not. Additional work by O'Hara [58] broadens the scope to include reading scenarios more generally associated with knowledge work, but in his work the distinction between reading by choice and reading by requirement persists. Each of these scenarios may involve one or more of the reading types described above (skimming, scanning, search, receptive or responsive reading).

Reading by choice

We will call reading by choice the category of reading types whose result is not subject to external evaluation and whose outcome has no bearing on the reader's academic or professional career. This type of reading is typically faster, has lower cognitive demand and higher affect (it is enjoyable, interesting or emotionally involving) [40]. The Lorch classification of reading by choice scenarios has categories that do not always seem distinct, so this classification has been augmented with longer descriptions.

We identify six scenarios that involve reading by choice [40] [58]:

- **Reading to apply** involves following a written plan or recipe in order to accomplish a specific task (e.g., reading a wedding magazine to plan a wedding, reading a DIY guide to get decorating ideas). It is linear, thorough and often interspersed with performing the task in question, involving short bursts of receptive reading.

- **Search** is reading to find text relevant to some goal or criterion (e.g., reading a TV schedule to decide on a program to watch, reading a sports page to find the score in a particular team's game). The search task may involve search reading [76] if the precise form of the goal is known or scanning [76] if it is unknown. Search is linear, cursory and does not necessarily result in comprehension of the document as a whole.

- **Reading to Self-Inform** is intended to fill a gap in the reader's personal (not professional) knowledge (e.g., reading a non-fiction book for information, reading a biography). It is one of the more cognitively demanding forms of reading by choice, although its outcomes are not evaluated. Receptive reading is common to this scenario, although responsive reading can also be employed.

- **Intellectually challenging reading**, in contrast to self-informing reading, has as a primary goal the engagement and exercise of the reader's intellect. Satisfaction is derived not from gaining new knowledge but rather from reaching outside one's comfort level in terms of concept and domain difficulty.

- **Reading for stimulation** is the type of reading most often associated with fiction (e.g., reading a mystery novel for suspense, reading a bestseller to escape). It aims to derive affective or emotional satisfaction from the text. Like watching a play or listening to a symphony, this type of reading is linear, receptive, focused and emotionally engaging, with few interruptions and an engaged imagination.

- **Light reading** (e.g., reading the comics in a dentist's office, reading junk mail) is also done for positive affect, but unlike reading for stimulation, it is superficial and non-engaging, and its primary purpose is relaxation rather than immersion.

Reading by requirement

Broadly, reading by requirement occurs when the result of the reading must be evaluated in its achievement of self-imposed or external goals and its outcomes matter to the reader in some definitive capacity.

These reading scenarios produce artifacts such as summaries, notes, or corrected versions of text, contribute to the reader's knowledge base, or are evaluated in the traditional sense in an academic or professional setting. These scenarios generally require deeper engagement with the text, more focus and less distraction.

There are nine scenarios that involve reading by requirement [40] [58]:

- **Reading to learn** applies to any type of reading that increases the reader's knowledge in a professional or academic setting without a specific evaluation looming in the immediate future. The learning in question may be procedural (e.g., how to solve a particular problem) or declarative (e.g., understanding a topic). Reading to learn involves multiple read-throughs at different levels of precision, self-testing and revision after the reading task.

- **Reading to memorize** text verbatim occurs when a pledge, script or quotation needs to be committed to memory. It is interspersed with mental repetition, frequent progress review and is extremely slow, involving a very high level of focused re-reading.

- **Test preparation** occurs when readers (typically students) are preparing for a time-limited written evaluation of their knowledge. It is a complex type of reading that is slow, meticulous and cognitively demanding. The implication of test preparation reading is that the original source documents or any additional content will be unavailable when the knowledge is tested, so while test preparation produces many artifacts such as external notes and summaries, its overall purpose is to internalize knowledge rather than to compile documents or memory aids.

- **Reading to summarize** is intended to extract the overall meaning and the important points from a document. It may or may not produce a summary as a final artifact, but it typically involves forms of annotation, parallel note-taking and an increased awareness of the structure of the text. This type of reading is often performed in academic settings and summarization is a companion task that has successfully been used as a proxy task for the evaluation of reading comprehension and learning [59].

- **Reading for research** is a type of reading that by nature involves managing multiple documents. This type of reading, ubiquitous in academia and common in knowledge-based workplaces, requires deriving a complete understanding of a topic through reading multiple documents and extracting, contrasting and evaluating information from each. The additional constraint of deciding *which* documents to include in this synthesis makes this activity extremely demanding in terms of cognitive load, mental representation and navigation.

- **Discussion preparation** for a class or meeting consists of a skimming read-through of the document(s) to be discussed with a focus on awareness of the main talking points of relevance to the discussion audience. Details are skipped over, with the focus being on the main points.

- **Proof-reading** is a low-level, high cognitive load reading scenario that focuses on correcting the spelling or grammar of a document. The reader is constantly interleaving the task with annotation. This reading type demands meticulous attention to text at the word and sentence level, but not necessarily an overall understanding of the major concepts.

- **Reading for revision** occurs when the overall structure and message of the document is within the reader's control and requires improvement. The reader has a complete mental model of the document as it is at the beginning of the task and a notion of what should be different about the document at the end of the revision process. This process is non-linear, it is demanding, it occurs at a higher level than proof-reading and it involves identifying passages for rewriting, restructuring concepts in the document outline and so on.

- **Reading for critical review** involves evaluating a document for validity and whether it furthers the knowledge area to which it contributes. This type of reading entails a thorough awareness of the domain in which the critiqued document resides, with more general recommendations than reading for revision. Verifying internal consistency of the argument presented involves non-linear navigation and requires that the reader develop a sense of the document's structure and the location of crucial ideas.

1.2.3 Reading Strategies

In the title of their seminal work on reading, van Doren and Adler [2] imply a crucial deficit in most readers' awareness of their metacognition while reading. The book, entitled "How to read a book", offers a four-level strategy that incorporates information triage, gathering important points, and considering the book within an ecosystem of knowledge, all activities that go well beyond the conversion of written symbols into language in the mind. Moreover, these levels and strategies involve conscious effort on the part of the reader and his or her recognition of the reasons for reading and mindfulness of the overall goals of the reading task. This characterization makes van Doren's levels consistent with the metacognitive theory of reading described in Section 1.1.4.

Sounding out sequences of characters, what we learn as reading in elementary school, is only the first and most basic level of reading. Then, a set of ad hoc or consciously acquired additional strategies must be applied to correctly represent what is being read and to incorporate it into the reader's existing body of knowledge.

Depending on the reading goals and situations, Waller [86] mentions "active" reading strategies that are notably different from the linear "sound-out" perception of reading described above. Multiple strategies and strategy bundles have been observed, taught and recommended to readers [27]. Some involve a modified engagement with the text: initial skimming, reinspection, and reciting headings and subheadings. Others require cognitive processing: posing questions about the text and rehearsing answers, and drawing inferences about the text. Others still result in the creation of other documents or artifacts such as summaries or study notes. While there are many methods and prescriptions available for readers to learn, individual levels of training and styles usually result in differences in strategy, process and awareness for different readers. However, many of them involve more complex reading patterns and companion activities such as non-linear navigation, annotation and parallel notetaking.

1.3 The Changing Nature of Reading

The previous sections dealt with the mind's representation of text. For the purposes of our discussion, text was defined as an unchanging linear document which was intended to be read from beginning to end, exclusively on paper. In stark contrast,in the last two decades we have increasingly begun to read digitally. Moreover, much of what we read no longer conforms to the properties of printed text. Static, page-by-page documents share our reading attention with scrolling webpages, document content and layout changes frequently and effortlessly, and documents often link to other documents.

As a result of the changing nature of reading materials, our reading strategies and habits are also changing. In a reading habits interview study, Liu [39] found increased browsing and scanning, keyword search, one-time reading, selective reading and non-linear reading. In contrast, sustained attention, in-depth reading and concentrated reading all decreased. Disturbingly, Liu also found that companion activities such as annotation were far more difficult digitally: over half of the participants interviewed indicated that they never annotated digital documents. Additionally, the multiple paths of non-linear navigation contribute to increased cognitive load on readers. These results are almost a decade old themselves, and cover the transition between paper and computer screen. They predate the widespread advent of tablets, pen-enabled tablet PCs and dedicated e-paper electronic readers. Further research will determine how these innovations contribute to the new balance of paper and digital reading.

It is worth noting that while e-readers represent the most obvious innovation in the reading ecosystem, they are essentially niche devices. They only support in-depth, linear reading, which is not the only or even the most widespread way we read. Despite technological advances across the digital reading domain, including faster PC's, better monitors, portable large-screen tablets and hybrid computers, the paperless office [69] has yet to come to fruition, seemingly on account of poor digital support for more complex reading tasks and companion activities.

As the knowledge economy grows in size and importance, the effects of digital reading on careers dependent on problem solving and information retrieval may be significant. Knowledge workers such as scientists, lawyers, engineers and architects spend 38% of their time retrieving information [51], mostly in written form, meaning that advances in the field of digital reading may have far-reaching, economically relevant and groundbreaking consequences. Similar effects could be expected in academia, where the strategies and skills of scholars are being formed in preparation for entering the knowledge work sector.

1.4 Responsive Reading

The type of reading alternatively described as responsive [76], analytical and inspectional [2], active [68] or close [36] reading is "the process of developing new knowledge or modifying existing knowledge by engaging with the ideas presented in a text." [76] It is characterized by the additional strategies and activities described in [86] and [27], and often includes annotation and parallel note-taking. The ability to perform this type of reading is a complex cognitive skill, which uses a great number of cognitive mechanisms to encode information, integrate it with previous knowledge, and retrieve it more easily as required. While "active reading" is a commonly used term, we agree with Thayer [76] that it implies that other forms of reading are non-active or passive in some meaningful way, which is not the case. Therefore, we will refer to this type of reading as **responsive reading** since this designation hints at the bidirectional nature of information flow in this type of reading, demanding the interweaving of content creation and content consumption.

1.4.1 The Requirements of Responsive Reading

As the activity in which knowledge workers most often engage, responsive reading has been widely studied, and a set of requirements for successful responsive reading has emerged. These requirements range from hardware specifications such as portability and basic text readability, through support for companion activities, to broad considerations such as archiving, information triage and multi-document navigation.

(Categories adapted from [16])

Hardware and physical requirements

- **Micromobility** [49], [45], [88] is the ability to reposition the device containing the document to obtain a more comfortable position or viewing angle. Such small adjustments are necessary to prevent fatigue.

- **Readability** [88], [93], [92] is the combination of typeface, font size, contrast and appearance that makes reading the document less tiring. These properties are rigidly determined for printed books but may be customizable for digital documents. Due to screen glare, backlighting or screen resolution, certain types of digital hardware are less readable than paper in the general case.

- **Support for Writing** [59] [1], [60], [69], [76] alongside reading is important, especially to responsive readers. This criterion covers both handwriting and typing and generally indicates the need to support written content creation while reading.

Page Level Requirements

- **Receptive (linear) reading** [63], [2] is the ability to read a document in a linear fashion, from start to finish. Thankfully, this requirement is the first one to be satisfied with all reading technologies.

- **Graphics and illustrations** [59], [49], [93], [11] are often used to enhance the understanding of, or to convey and interpret crucial information in conjunction with, the text. [38] A responsive reading platform must support graphics in as much detail as is required. Many academic fields (e.g., biology, physics, geography) rely heavily on rich and detailed graphics in their written documents, and poor support for this criterion due to low resolution, colour display issues or poor reproduction quality can result in documents being virtually unintelligible.

- **Superimposed freehand annotations** [2], [43], [76] are a crucial component of the responsive reading process and many studies of digital readers cite the lack of support for such annotations as the primary reason for non-adoption of new technologies.

- **Text search** [88],[93] can greatly facilitate the searching and scanning types of reading, and is a requirement for which any digital representation of text has a tremendous advantage over paper.

- **Glancing back to re-read** [47], [48] a recently read passage is necessary to recover focus in the event of an interruption and to follow complex arguments in text. Most reading technologies intrinsically support this requirement. In the worst case, on small-screen devices like smartphones and smart watches, glancing back may require an explicit navigation operation.

- **Physical/Kinesthetic cues** [59], [76] are vestigial to paper documents. They are the unconscious indications the length of a document, how far into a document the reader has progressed, or the immediate haptic clue the reader receives when he or she lifts a page in order to turn it. More broadly, these are the result of the fact that when a paper document is read, its length and the act of advancing through it affect its physical configuration. In contrast, the tactile experience of digital reading is not noticeably different based on document length or position. These clues are useful in the reader's orientation in the document.

Document Level Requirements

- **Non-sequential navigation** [59], [11], [74] is crucially important in responsive reading, where related concepts may be mentioned pages or even documents apart. The ability to quickly flip back to a table of contents or forward to a list of references may be more difficult to support on devices with a slow refresh rate and page-based sequential navigation.
- **Cognitive maps** [59], [88], [76] are spatial or conceptual representations of a text in the reader's mind that facilitate navigation and awareness. They are overviews of the document that are obtained through a single document-wide view or through an initial receptive reading.
- **Skimming to get overview** [63], [2] is an initial evaluation requirement that involves rapid linear navigation. Devices that require frequent pageturns due to small screen size and those with slow refresh rates would make this process more difficult.
- **Discovery of topical knowledge** [63], [2] is similar to skimming, but has as a goal extracting information not about the document as a whole, but about a specific concept or subsection of the document.
- **Switching between navigation styles** [76] according to the reader's needs is easy and effortless on paper. Readers may turn a page, fold a spine in half or use a finger as a bookmark in a previous section. Having multiple ways of navigating through a document that are equally easy to activate contributes to maintaining engagement with the text and navigating without interruption.

Workspace Level Requirements

These requirements concern the interaction between, and management of, multiple documents in a single workspace.

- **Sorting and triage of documents** [2], [50] is an important skill to knowledge workers. Deciding which documents are crucial, which are helpful, and which are unrelated to a research topic is a skill that must be, and is demonstrated by every Ph.D. hopeful on the planet. Technology may help or hinder this requirement depending on the way it manages multiple documents.
- **Spatial layout** [59], [60] is the ability to freely arrange multiple documents in sensemaking ways.
- **Extracting information** [1], [56] such as important snippets or data is a writing activity that frequently accompanies responsive reading. On paper, this entails copying or photocopying information, while digital text often supports copy-paste functionality.
- **Reading from multiple documents** [59], [60], [88], [74]) for comparison or contrast is another requirement that is easily supported by a physical desktop and can be challenging in the digital realm due to small screen size or layout rigidity.

- **Integrating with PC workflows** [76], [74], [48], [60], [56] is a crucial requirement in knowledge work. Most knowledge workers use a PC as their primary digital device, often to create and disseminate documents related to their reading. Regardless of the reading technology, the transition of some version of the document into a PC workflow is often inevitable and must be supported.

Multi-session Reading Requirements

- **Reading in different venues** [69], [76], [74], also known as macromobility, refers to the ability to read a document in more than one location and indicates the general portability of a reading technology. Macromobility may be necessary when commuting, bringing in documents from another source (e.g., a library) or sharing documents with others. A 26-volume encyclopaedia is largely non-portable in its paper incarnation while its online equivalent may easily fit on a portable reading device.

- **Filing and archiving** [41], [69] involves the proper storage of documents once they have been processed so that they may be preserved, as well as retrieved or consulted at a later date. Most knowledge workers are not explicitly trained in filing, meaning many have piles of printed papers on their desks or flat digital folders full of PDF documents on their hard drives without any storage or filing system. The ability to organize and archive is crucial to avoid misplacing or forgetting documents.

- **Restoring reading workspace, resuming reading activities** [41], [88], [74] is often sought in service of the visuospatial awareness readers have built up about their workspaces and the documents within them. On paper, it is simple, yet inconvenient, and involves leaving the documents on one's desk and locking the office door. Digitally, much more may be possible, such as saving multiple virtual desktops and cycling through them depending on the task at hand.

Chapter 2

Reading Technologies

In the millennia since the first use of paper for writing, it has remained the dominant reading technology in the world, largely due to the lack of alternatives. Whether the writing surface was made from pressed papyrus, vellum or tree pulp depended only on the skills of the respective tradesmen and the availability of said commodities. Even the shift from the scroll to the bound pages of the codex we are familiar with today did not intrinsically change the nature of reading as an activity. Paper has important advantages such as familiarity, materiality and persistence, but also comes with production and storage costs, has a non-negligible mass and a greatly restricted set of interaction affordances.

It is therefore significant, if unsurprising, that the past one hundred years, which have seen exponential advancements in population, connectivity and technology, have also generated more alternatives to paper for reading than the previous six thousand. A trend begun with the personal computer, or, more accurately, the personal computer monitor, currently has almost a dozen different categories of devices that could conceivably be used for reading. The desktop PC has served as such for over a generation, with its portable counterpart the laptop appearing not long afterwards. Both are extensively used in education and knowledge work, yet some of their properties contribute to unresolved issues in digital reading. More recently, e-paper reading devices are attempting to mimic the optical properties of paper to provide a more familiar and enjoyable reading experience, while tablets and smartphones offer tremendous storage capacity and processing power in portable packages of varying sizes. The latest generation of devices involves a hybridization of the laptop computer and the tablet, with devices including detachable screens or keyboards. Interestingly, although they have been available for decades, stylus devices that enable freeform writing and drawing have failed to gain significant traction in the mainstream, even though freehand annotation is one of the most cited issues with digital reading. Some attention has also been paid to the use of large displays as reading workspaces, specifically in tabletop computers, which could conceivably contain an entire virtual desktop.

Alongside mainstream reading devices, dozens of research prototypes and niche products have attempted to bridge the gap between paper and digital, offering anything from flexible screens that could be turned like paper pages to multiple pen-enabled slates to help with document management, but these prototypes have often taken an insular, competitive and comparative stance against all alternatives. Demonstrating a slight increase in reading comprehension in a laboratory study or favourable user attitudes is often uncorrelated with the likelihood that the device will reach production, or that users will in fact adopt it over its competitors.

2.1 Paper

Before we explore the wide variety of devices that enable digital reading, let us consider the original reading medium: paper. Usually white or light in colour with dark symbolic markings, paper requires an external light source to be read, but it is typically matte, exhibiting high readability and low glare. Paper may persist for centuries if stored correctly, it is immune to format or medium obsolescence and it requires very little technological overhead to be read. A printed document contains text that is sealed and finalized at the moment of printing. Once produced, its layout and the information contained within remain constant. Its navigation techniques are simple, familiar and habitual for all literate humans. Paper is flexible, foldable and accepts printer, ink and pencil markings. Paper requires no additional infrastructure to be used in educational settings.

Paper is material, meaning that each document is a distinct object that can be manipulated separately. Books and papers may be arranged arbitrarily on a working surface. The materiality of paper is a comfort to less technologically inclined readers, but it also poses significant challenges. Alongside the cost and environmental impact of manufacturing paper, there are significant per-unit costs to print, bind, and distribute paper books. Availability and acquisition is also material, meaning that a book may not be in stock where the reader would like to purchase it, it may only be available after a significant delay and costly shipping from another location, or it may be out of print. In the case of some documents, the availability of on-demand printing makes it possible to acquire a document digitally, then print it on paper. This convenience perhaps contributes to the persistence of paper-based reading in the office, despite long-standing predictions to the contrary [69]. However, due to the linear relationship between document length and printing cost, on-demand local printing is only feasible for shorter documents and most books are still subject to the material procurement limitations described above if they are to be read on paper.

There is another, curatorial issue with paper books, and that is the tremendous investment of time and effort required to publish a book. Paper-based publishers have to deem a book worthy of printing or reprinting based on nebulous and opaque criteria in order for the book object to be manufactured and made available to the reader. As with technology, mainstream appeal is a principal driving factor behind these decisions. Finally, paper books must be stored, and while a single book is relatively portable, their non-negligible weight makes it impossible for readers to carry more than a few with them at any one time.

Regardless of its suitability to responsive reading, digital format text is staggeringly less expensive to produce, store, and much faster to replicate and disseminate. A reader may carry a library containing thousands of volumes in a pocket-sized storage device, or she may opt to access and instantly download books she requires through an Internet connection. An update or a new edition of a digital text results in the updating, not the discarding, of older editions, and digital versions of books remain pristine and unspoilt regardless of the number of times they have been flipped through. In addition, while a print book has at most a handful of editions to choose from, severely limiting the availability of different font sizes and typefaces, a digital version may be customized to the reader's exact preferece.

2.2 Mainstream Devices for Digital Reading

There are several types of mainstream devices on which digital reading is possible and frequently occurs. In this section, we will define each one and determine how well they support each of the reading requirements we have identified in Chapter 1. It is important to note that, for reasons to be detailed in Chapter 4, our analysis of support for responsive reading will focus primarily on existing mainstream technology (i.e. mass-produced devices available to the general public). While there is a tremendous variety of prototypes and augmentations that stretch the possibilities of digital reading, none of them are likely to factor into the current workflows of the vast majority of responsive readers.

The most visually significant hardware aspect of a digital reading device is its screen. Screen size is a salient property, and mainstream reading devices range in diagonal screen size from 3 to 50+ inches. However, other qualities of screens such as colour saturation, contrast, screen resolution and refresh rate are also significant to reading. Screens that have been or are used for digital reading rely on a variety of technologies, each with its own strengths and limitations.

Cathode ray tube (CRT) displays were used for reading when there were no other options in mainframe terminals or personal computers, but the weight-to-size ratio of CRT technology and the fairly significant limitations in resolution make them largely unsuitable for digital reading, and entirely unsuitable for embedding into modern reading-capable digital devices. As an outdated exception, CRT monitors are still used in high-magnification closed-circuit television systems for the visually impaired, although they typically magnify paper documents and non-digital objects such as medication bottles and product labels.

The liquid crystal display (LCD) was the first flat, portable display option available for digital devices, and such displays are found on many laptops and early dedicated e-readers. While resolution is better than that of CRT's and sufficient for most reading tasks, LCD screens tend to exhibit narrow viewing angles, high glare and muted colours. In addition, they are subject to visual distortion if flexed or pressed, making stylus use uncommon and undesirable.

Organic light-emitting diode (OLED) displays are the current standard for televisions, computer monitors, smartphones and tablets. With bright colours and increasingly high resolution, these screens suffer only from a typically glossy finish that invites glare, requiring greater screen brightness, which reduces battery life and contributes to reader fatigue.

E-paper or electronic ink displays have some notable properties that make them more suitable to digital reading and less suitable for general purpose computing. E-paper screens have a limited colour palette (typically restricted to 16-shade grayscale) and a far lower screen refresh rate due to the electromagnetic properties of their pixels. Conversely, since pixel polarization is persistent, e-paper screens can maintain an image on display without using battery power. They approximate printed pages in contrast and visual quality and, unlike other display technologies, they require an external light source to be read. While this technology has been applied globally to dedicated e-readers and haltingly to wristwatches and smartphones, its colour and resolution deficiencies make it poorly suited to rendering graphs, figures or illustrations.

Although screen type and quality play an obvious part in the suitability of a digital device to the task of reading, considerations of how devices support the requirements of responsive reading, and additional features such as connectivity, general-purpose usability, configurability and cost must also be taken into account.

2.2.1 Definitions

We will differentiate between several types of hardware without claiming that these types are mutually exclusive. Over the last several years, the versatility of hardware configurations in portable devices has allowed for many devices to straddle multiple categories. For instance, a Microsoft Surface Pro can be used as a laptop while placed on a desk and attached to a keyboard dock, or it may be handheld like a stylus tablet. Each of these modes has different implications for the act of reading, so they shall be considered separately, even if they employ the same device.

The devices used for digital reading are:

- **Desktop PC**: Connected to a number of external peripherals, including one or more monitors.
- **Laptop PC**: A portable computer that incorporates a monitor, keyboard and pointing device.
- **Stylus PC**: Often used in tablet form, these PC's are meant mostly for pen input and notetaking.
- **Tabletop**: Large, horizontal, table-like surfaces meant to display a full digital workplace.
- **Tablet**: Touch-based slate devices, largely for content consumption.
- **Smartphone**: Small, touch-enabled devices used primarily for communication.
- **E-reader**: Reading-focused tablets featuring low power, high-contrast greyscale displays with a slow refresh rate. *

* N.B.: Some e-readers like the Kindle Fire are equipped with OLED displays, making them more akin to tablets in the properties of their screens, despite the fact that they are reading-specific devices.

Device name	**Screen**	**Keyboard**	**Mouse/Trackpad**	**Pen**	**Touch**	**Portable**	**Examples**
Desktop PC	14-25" CRT/LCD/OLED	●	●	○	○	○	Mac Pro
Laptop PC	10-17" LCD/OLED	●	●	○	○	●	Lenovo ThinkPad
Stylus PC	9-14" LCD/OLED	◑	◑	●	○	●	ThinkPad Tablet 2
Tabletop	40" LCD	◑	◑	◑	●	○	Microsoft PixelSense
Tablet	7-10" OLED	○	○	○	●	●	iPad, Nexus 7
Smartphone	3-4" OLED	○	○	○	●	●	iPhone
E-reader	4-7" e-paper	◑	○	○	◑	●	Amazon Kindle

Figure 2.1: Types of reading technologies

2.2.2 Support for Responsive Reading Requirements

(Adapted and expanded from [16])

LEGEND
○= Poor support
◑= Moderate support
●= Excellent support
(●) = 'Target' technology

	Desktop PC	Laptop PC	Stylus PC	Tabletop	Tablet	E-reader	Smartphone	Loose Paper	Bound Paper
Hardware/Physical Requirements									
Micromobility: thin, light, graspable	○	○	◑	○	●	●	●	(●)	◑
Readability	◑	◑	◑	○	◑	●	◑	●	(●)
Support for handwriting	○	○	●	◑	○	○	○	(●)	◑
Support for typing	(●)	●	○	◑	◑	○	◑	○	○
Page Level Requirements									
Receptive (linear) reading	●	●	●	●	●	●	●	●	(●)
Illustrations	●	●	●	●	●	○	◑	◑	(●)
Superimposed annotations	○	○	●	◑	○	○	○	(●)	◑
Text search	●	●	●	●	●	●	●	○	○
Within-page rereading	◑	◑	●	●	●	●	◑	(●)	●
Kinesthetic cues	○	○	○	○	○	○	○	●	(●)
Document Level Requirements									
Non-sequential navigation	●	●	●	●	●	◑	●	(●)	●
Cognitive map of content	◑	◑	◑	◑	◑	◑	○	◑	◑
Skimming to get overview	◑	◑	◑	◑	◑	○	○	●	(●)
Discovery of topical knowledge	◑	◑	◑	◑	◑	○	○	●	(●)
Switch between reading styles	◑	◑	◑	○	○	○	○	(●)	●
Workspace Level Requirements									
Sorting documents	◑	◑	◑	●	○	○	○	(●)	◑
Spatial layout	○	○	○	●	○	○	○	(●)	◑
Extracting information (copy-paste)	(●)	●	●	●	◑	○	○	○	○
Reading from multiple documents	◑	◑	◑	●	○	○	○	(●)	◑
Integrating with PC workflows	(●)	●	◑	◑	○	○	○	○	○
Multi-session Reading Requirements									
Reading in different venues	○	●	●	○	●	●	●	◑	◑
Filing and archiving	(●)	●	●	●	◑	○	◑	◑	●
Recreating reading workspace	◑	◑	◑	◑	◑	◑	○	○	○

Figure 2.2: Support for responsive reading by technology

The table above is a simplification of multiple considerations and facets at the intersection of each category and technology. Each cell merits its own research and full exploration, but even without undertaking such an endeavour, we can generalize about the ability of any reading technology to support the full range of responsive reading requirements. To interpret the table, we consider the number of completely and partially filled circles along each column, representing a particular reading technology. Based on a list of requirements that is naturally biased towards paper, it is unsurprising that paper remains the target technology in many categories. While that is the case, no attempt has been made in the table to rank requirements against each other, and some of paper's failings such as text search, typing and copy-paste support may be more significant than they appear. The table also excludes considerations that do not immediately pertain to reading (e.g., battery life, support for other important activities, availability of content, etc.).

The most important thing to note is that, among these device categories that represent the complete range of consumer-level digital technology suitable to reading, there is no clear winner: no single technology that seems to have emerged as the natural successor to paper. The stylus PC, which offers some of the best support for freeform annotation, is the closest there is, but it is ironically among the least popular mainstream devices in the table, and its support for multiple document views is unremarkably poor.

Naturally, these devices can be combined and augmented with hardware or software to better support one or more of the requirements above, but it seems that even in the research literature, prototypes have failed to produce a universal responsive reading device. The following section is a partial summary of some of these efforts, noting relevant achievements and identifying general trends.

2.3 Related Research and Prototype Technology

As with any technical field of research, there is a large number of prototypes in digital reading that have been built to explore or address a specific shortcoming or aspect of the domain and, having done so, never entered mass production. Some devices were produced and marketed to some excitement only to be abandoned by the market months after launch. Predictably, many prototypes have been geared towards bridging the gap between the digital and paper-based experience, recognizing the advantages of paper such as familiarity, readability, cost and support for annotation.

The requirements of document reading and annotation led to the creation of the iLiad e-reader in 2006. It featured an 8-inch or a 10-inch e-paper display, a stylus and digitizer, and multiple interfaces for loading documents onto the device. The device failed and the company filed for bankruptcy in 2010, although it had many features that have been absent in subsequent commercial devices: a stylus, an extensible Lunix-based OS, and a wide range of supported file types that could be loaded from a variety of storage media.

Another feature of the iLiad was that its larger version had the largest screen of any e-book reader available at the time. Further exploring the issue of screen space, the Skiff reader was an 11.5-inch flexible e-paper device that could display an entire Letter-sized page without shrinking it, which remains a compelling distinction among e-readers today. Its prototype was announced in 2010 but the device was abandoned, never entering mass production. Regardless, it was one of the first devices to feature a thin, flexible e-paper display that could bend over 90 degrees.

The notion of screen flexibility has been explored since the creation of e-paper, while numerous efforts have tried to recreate the kinaesthetic navigation action of page-turning. As an example, Tajika et al. [72] introduced a flexible plastic sheet on top of a LCD display that could trigger page-turning or page-bending actions based on how it was manipulated.

As flexible display technology improves, so do the opportunities to match other properties of paper. In their latest incarnation, HD colour flexible screens were demonstrated at the 2013 Samsung keynote at the Consumer Electronics Show. Samsung's primary emphasis was on the durability and versatility of such displays. The technology was featured in a wraparound smartphone display, but concept prototypes of a foldable tablet and a rolled-up display were also presented. The former is eerily reminiscent of a bound codex, while the latter behaves like a scroll. Despite these advancements, whether the flexibility of the paper page and its turning affordance is significant enough to be emulated in digital reading devices remains to be seen.

In stark contrast to the push of technological advancements, some researchers have found the benefits of digital reading such as portability, storage capacity and low per-document cost to be so significant that they have explored adapting and using low-cost digital devices that were never meant to serve as e-readers in developing world classrooms. Thinyane and Thorne [77] used digital picture frames as makeshift e-readers in a South African classroom to moderate success. Despite lacking almost all the features required for responsive reading, the picture frames still brought noticeable improvements in document access in those classrooms.

In leisure reading, the familiarity of paper and the limitations of e-paper displays have made the transition to electronic reading slower than anticipated, especially for children's or illustrated literature. This has pushed e-reader application developers to create enhanced e-books that not only match and mirror the aesthetic qualitiesof traditional books, but also add animations, sound and special effects, creating a multimedia experience that stretches the definition of reading. A version of Arthur Conan Doyle's "Sherlock Holmes: The Adventure of the Speckled Band" created by Byook and marketed as an "augmented e-book" presents faded and worn digital pages with non-optional background animations and sounds that can be activated with a touch. While this form of enhancement is a definite step in bringing printed text to life, it is unclear how visual distractions affect the task of focused reading and how such enhanced e-books fare against their static counterparts in comprehension and reading skill acquisition.

Freeform ink annotation is a crucially important element of responsive reading and while the bulk of commercially available devices today does not support it, the research community has long attempted to enhance e-readers with the ability not only to annotate as seamlessly as one does on paper, but also to save, transfer, synchronize and summarize annotations. XLibris [68] was an e-paper digital reading appliance designed in 1998 to address some of the issues of digital responsive reading with support for freeform digital ink annotations. In a reading group deployment, readers generated as many annotations with XLibris, and as easily, as on paper [47]. Furthermore, XLibris enabled readers to look at the text snippets they had annotated in an aggregate view called the "Reader's Notebook". This technology hinted at future developments that would enable automatic summarization of text based on highlighted or annotated passages, a useful feature that paper cannot provide.

Finally, several research prototypes have attempted to emulate the multiple document display capabilities of paper, as well as the ability to spatially arrange paper artifacts. Chen et al. [15] began by prototyping a dual-display 10" stylus-enabled e-paper digital reader geared towards responsive reading. The two displays were arranged side by side in a portrait orientation reminiscent of adjacent pages in an open book. While they could be attached together to display two consecutive pages from a digital document, they could also be used and manipulated separately. Either device could display a blank page for notetaking, a table of contents, or a navigation-facilitating overview of the document.

An extension to the prototype imagined a workspace in which multiple homogeneous devices were available to the reader. Extending the number of e-paper stylus-enabled devices from two to four, Chen et al. created the United Slates framework [16], which enabled devices to be arbitrarily arranged across a workspace and to display different pages, documents, or views. Better yet, they could display and instantly synchronize the same page across multiple devices, enabling collaborative work. Evaluation with graduate students [17] elicited positive feedback, especially for the devices' annotation capabilities, seamless PC integration, intuitive document loading, and synchronization.

Although devices such as Chen's e-reader could become inexpensive enough that his multi-device ecosystem would be a real possibility, it would pose difficulties in migrating from one location to another. Furthermore, it is unclear if users would be willing to own and adopt multiple homogenous devices solely for the purpose of responsive reading.

A plausible scenario resulting from the current state of technology ownership among students advocates for the use of heterogeneous devices in the service of responsive reading. If a reader already owns a laptop, a tablet and a smartphone, and has access to a tabletop display or a large TV on occasion, a logical step would be to attempt to integrate these devices into one reading workflow. Arthur et al. [5] developed a toolkit for display annexation: beaming content to any display in the vicinity. Further work currently under review by Yang and Wigdor [91] enables defining arbitrary views and sending them to various devices, bypassing OS and platform differences, while Hamilton and Wigdor [28] are applying novel interaction techniques to the task of synchronizing and managing data across devices. Although not yet adopted into mainstream use, such multi-device work signals an important potential shift away from the prevailing trend of closed devices with manufacturer ecosystem biases in which data sharing is far from seamless and screen sharing among heterogeneous devices is practically impossible.

Chapter 3

Companion Activities to Reading

Responsive reading does not occur in isolation. Multiple documents are often consulted, and documents are annotated, summarized, skimmed and arranged in meaningful ways. There is a set of essential companion activities to responsive reading, and many of them such as freeform annotation, multi-document navigation and random access are far better supported on paper than digitally. Other activities such as indexing, clipping and hyperlinking are only possible in a digital format. Many of these activities produce additional content or alter the source material in some way, and they must be managed alongside the document.

The most significant way in which responsive reading differs from casual or receptive reading is through the meaningful use of annotation to process, comment on, and interpret the document. Annotation occurs unselfconsciously, directly on the document and may include highlighting, underlining, marginal notes and a huge variety of symbols, colours and codes that vary from reader to reader. At the moment, this activity is effortlessly supported on paper, but requires mode switching or training on all but a few kinds of digital devices.

Another difference in responsive reading is the complexity added in navigating documents. A receptive reader navigates almost exclusively by turning individual pages in the course of his or her reading. In contrast, a responsive reader may flip ahead, skim, re-read a passage or consult a reference section multiple times in a single reading session. More importantly, due to the demands of critical thinking and the nature of knowledge work, responsive readers are often reading from multiple related documents and navigating rapidly between them. A single device's screen, no matter how big, may be insufficient to cover the full range of necessary arrangements and operations required with multiple documents.

This places heavy demands on the reader's ability to keep his or her orientation (i.e., place) in each document. Documents can be meaningfully organized across a workspace, and superficial features such as graphic placement or page boundaries may help the reader quickly locate and navigate to passages of interest. Spatial workspace arrangement and session preservation facilitate inter-document navigation. Other activities and considerations must also be supported, including notetaking, clipping, and content sharing.

Despite paper being the standard on many of these essential activities, migrating to digital formats has two notable advantages: documents need not be static, and they need not be linear. An online document can be updated as often as is necessary, multiple times a day, to contain the most relevant and up-to-date information. In addition, a document can be structured to take advantage of the non-linear affordances of hypertext: blocks of text connected by deliberate links. In fact, every webpage in existence leverages the reader's ability to take different paths through the content, finding what he or she needs by following relevant links. Non-linearity and hypertext are changing reading patterns and online hyperlinked content represents an increasing proportion of all reading materials, but the support for companion activities in Web browsers is sparse and the potential of non-linear educational content is still not fully explored.

3.1 Annotation

For the purposes of this document, we will treat **annotation** as the process of making informal markings on the document itself in the course of reading, a way of engaging with the text through ad-hoc content creation such as highlighting a passage, adding a note in the margins or linking two passages with an arrow. Annotations vary widely in purpose, content and desired permanence, but they all share a few common properties, namely that their creation is interleaved with the act of reading, they pertain to a specific region in the text and that, on paper, the practice of annotating is seamless, flexible and well-developed. Unlike reading, annotation is a content creation, rather than a content consumption activity, meaning that it is subject to the same constraints as other digital input techniques, including text entry, handwriting recognition, mode switching, etc. Annotation has been poorly supported in the realm of electronic reading for several reasons, the main one being that the vast majority of e-books being purchased are titles for leisure reading. As a consequence all widely available dedicated e-readers and e-reading applications lack pen input support, being almost exclusively geared towards content consumption rather than creation.

3.1.1 Function in Responsive Reading

When reading responsively, the reader engages critically with the text for the purposes of altering his or her existing knowledge based on the contents of the document. This engagement often leads the reader to annotate the document being read. In fact, annotations are one of the main ways in which readers process and engage with text under responsive reading conditions.

Numerous studies [76] [59] [46] cite annotation as an ubiquitous companion activity to responsive reading. In fact, Adler et al. [1] reported that in professional environments reading occurs in conjunction with writing more often than in isolation, and that document annotation accounted for over a quarter of time spent writing while reading. With the exception of form filling (which is an essential activity to workers but has no bearing on responsive reading), annotation was the most frequent companion writing activity reported in [1].

O'Hara and Sellen [59] conducted interviews with ten knowledge workers who were asked to read and summarize an article on paper and one on a computer screen. Their participants engaged in constant interleaving of reading and writing, and greatly valued the seamless and effortless nature of annotating on paper while disparaging the context and mode switching required to do the same with the digital document.

Using annotations has proven benefits to learning outcomes in an educational setting. Wallen et al. [85] explored how the availability of different kinds of pre-authored annotations in a digital environment affected learning outcomes in college students and found that their presence correlated strongly with an increase in performance on learning outcome measures such as concept recognition, comprehension and domain transfer. This finding applied to preexisting annotations published alongside the source text, not ones authored by the readers themselves.

Wolfe [89] conducted a study on student-created persuasive essays that included annotations and found that reading the author's annotations may greatly influence a reader's attitude towards the text, improve recall of important passages and increase engagement with the material. Annotations, Wolfe noted, also served as a conduit for interaction between the reader and the annotator, influencing the reader's perception of the writer's competence.

In a study of annotation and underlining [29] on science and history texts for undergraduate students, two groups were trained in marginal notetaking and underlining (both forms of annotation), respectively and asked to use these methods while reading a science and a history text and complete a written retelling of both texts immediately after. One group was not trained in either method but read the same texts and was asked to complete an immediate retelling, and a control group was not trained and not asked to retell the passages. All groups completed knowledge pre- and post-tests on the two domains selected. The three groups that were asked to retell the passages immediately after, effectively generating additional content through annotation and summarization, performed better than the control group.

Interestingly, readers in the third group who were free to use their own annotation methods rather than the one prescribed in the training conditions for the first two groups performed better than either of the others, utilizing previously learned annotation strategies that they were free to apply as appropriate to their learning process, and better than the control group, which was under no obligation to read responsively. This result indicates that self-generated annotations unconstrained by an externally imposed strategy lead to improved learning outcomes, more so than applying an externally imposed dedicated annotation strategy.

The findings of [29] hint at the difference between an impositional approach to responsive reading, in which techniques and methods are suggested and imposed by a central authority, and an emergent approach, in which individual readers decide on the best strategies for themselves. As the study above clearly illustrates, the emergent cases often yield better performance, because they are able to adapt to individual differences in responsive reading habits, strategies and practices. In contrast, the impositional approach which large institutions often undertake in reading technology deployments often results in poor adoption and ultimate failure. The contrast between emergence and imposition, and the value of considering and exploring emergent techniques, are further discussed in Chapter 4.

The ability to annotate documents is reported as an essential and often requested feature of digital reading appliances for education. In Thayer et al.'s [76] semester-long e-reader deployment with graduate students, the consequences of offering no annotation support in a digital reader were dire: of the students who always annotated text, many abandoned the reader in favour of paper because annotation was too cumbersome on the electronic device. Two other university deployments by Young [93] and Behler [11] report similar results.

The majority of reading annotations are created for private, individual consumption, with no intention of ever being shared or made public. This grants the reader a certain amount of freedom, making the annotated document an extension of his or her thought process and interpretation. Such an interaction is by nature unselfconscious and informal, and the reader is only constrained by whatever system, organizational scheme or set of rules he or she deems useful.

According to Marshall [43], such private annotations serve one of the following purposes:

- Signal importance for future attention

- Aid recall of text
- Provide a space to work out a problem or equation
- Record interpretation or processing
- Focus reader's attention
- Reflect circumstances of the reading session orthogonal to the text being read

Some annotations may also serve a social, interpersonal or collaborative purpose. These annotations are created to enhance the original text for the benefit of a third party. For example, an instructor annotates a student-submitted essay as she marks it in order for the student to be able to trace her evaluation process. Copy editors and supervisors make notes and annotate documents so they may be revised by their author. Finally, collaborative documents may contain public annotations from many authors.

Marshall and Brush [46] made the distinction between annotations meant only for personal use and those that would eventually be shared with a group of readers. They discovered significant differences in the nature of personal and public annotations, namely that only a small portion of personal annotations were deemed by their authors useful enough to make public, and that personal annotations were often heavily edited before being shared. This conclusion hints at an additional step which complicates the workflow for annotations meant for public consumption.

Interestingly, Marshall [46] noted that although annotation is an invaluable companion activity to reading, the objective value of the annotations themselves was far lower than even their authors imagined. Many annotators who returned to a document several weeks later were disappointed by how little insight their annotations had to offer them at the time of a reread.

3.1.2 Components and Attributes

Marshall [44] classifies and explores the range of annotations that occur on paper and that should be supported in an electronic equivalent.

An annotation has three principal components: a *body*, an *anchor* and a *marker* [44].

Body

The *body* is the additional content generated by the annotation. It may be explicit (e.g., a textual margin note), or implicit (e.g. an underlined sentence). An annotation with an implicit body does not express the reader's interpretation as added text, but rather signals the importance of some portion of the original document.

Anchor

The *anchor* is the specific range of the text that the annotation applies to. Anchors may encompass the entire document or pinpoint a single word or even character. They may also be explicit (e.g. a pair of parentheses or the start and end of a highlighter section) or implicit (e.g. an asterisk in the margins that is only vaguely linked to a section of text).

Marker

The third component, the *marker*, refers to the visual characteristics of the annotation, which are necessarily distinct from the characteristics of the source text and may contain information themselves. A red underline and a yellow highlighter mark could span the same range and have the same (implicit) body but they may encode different information to the reader. In the digital realm, it may also be possible for the marker to purposefully deflect, rather than attract, attention from a passage, for instance by decreasing the contrast or brightness of unimportant parts of the text [74].

A study of the relative frequency of different types of annotations in used university textbooks [46] found that by far the most common type of annotation (83%) was underlining and highlighting which has an explicit anchor but no content. Content without anchors (e.g., a note in the margin not linked to any particular passage) accounted for 7% of annotations, while compound annotations made up 9% of the total. Finally, 1% of annotations had neither, meaning they were markings entirely unrelated to the text (i.e. doodles).

ANNOTATION	Anchor	No Anchor
Content	9%	7%
No content	83%	1%

Author

Public annotations have an additional property that should be considered: their *author*. As more people edit the same document, it becomes necessary to store each annotation's author alongside it, enabling filtering by authors and appropriate attribution.

Miscellaneous

Many annotations have a valuable purpose upon re-reading a document, but some do not. For example, the highlighting of a difficult passage that occurred while the reader was tracking her reading progress becomes obsolete almost immediately after its creation and the author may benefit from the ability to erase it or prevent it from being displayed before reading that passage again. In addition, annotations have different values at different times, so possible digital attributes could include an *expiration period* or *tags* that filter annotations by their intended purpose, priority or category [74]. Finally, storing the *timestamp* of an annotation's creation as an attribute can be very useful to tracing the annotation history of a responsive reader as she works on a document.

This range of attributes is not only possible to implement in a digital environment, but it makes digital annotations more richly represented than their paper counterparts, despite the considerable difficulties in enabling seamless and unselfconscious annotation authoring on digital documents. Storing more annotation attributes will enable digital reading appliances to accomplish tasks such as filtering, time-based review and deletion that are impossible on paper.

3.1.3 Metaphysical Role

A crucial quality of printed text is its propensity for wear and tear. There is a very significant difference between a brand new textbook that has never been read and the same textbook after a few different owners. Philosophers such as Baudrillard [9] deem this quality extremely important to the tactile and sensory experience of reading. The fact that a book becomes worn as it is being handled leaves a trace of the reading experience and grounds it in reality. In contrast, digital media such as CD's and electronic text, which are unaffected by being read, are impervious to our intractions with them ("It's as though you'd never used it. It's as though you didn't exist." [9])

Like standard wear-and-tear, annotations are signs of use, but unlike loose pages, dirty bindings or missing covers, which are implicit and incidental, annotations are explicit and purposeful. Annotations are reader-generated content, blurring the line between reading and writing and contributing to the engagement with the text at a tactile level.

Wear-and-tear, despite being part of the metaphysical evidence that we have read, is not generally considered a desirable outcome. Owners of printed books are often hesitant to mark up their copies, especially in pen, because they would be contributing to their book losing its canonical, as-printed, look. On the other hand, marking up a digital copy carries less finality, given that an unaltered version is always available, leading to less guilt in annotation [71]. Undergraduate readers in Chen et al.'s [15] dual display e-reader study reported writing more on their pen-enabled device than they would have on paper because they felt they were not *"defacing"* the text.

In O'Hara et al. [59], participants reported a high degree of discomfort when annotating a digital document because their markings were indistinguishable from the original text (which was presented in a word processor), indicating that readers would like to regard annotations as a separate layer that cannot be confused with or affect the original document. Since this behaviour was observed in a digital copy, the discomfort seems to be with altering the document's contents, rather than merely with damaging its material manifestation as would be the case in marking up a printed book. It appears that readers would like to make annotations quickly and without breaking their reading flow, but they would also like them to be decidedly distinct from, and to make no alterations to, the source document.

3.1.4 Digital Annotations Wishlist

A crucial reason why knowledge workers still print out and read text on paper is on account of its seamless support for interleaving annotation with the act of reading. Most digital reading appliances are currently geared towards casual or leisure reading, which is understood as a mostly linear content consumption activity. Annotations pose a content creation challenge which consumption-focused devices are poorly equipped to meet.

The following design guidelines are essential to supporting digital annotations [43], [59] [74].

- Annotations should be *in situ*, on the document itself, not in a separate view [44], but they should also be distinguishable from the original text [59].

- Annotations without explicit content such as highlights and underlines should be supported, given that they make up the vast majority of markup [46].

- Annotations should be freeform and fluid. Given that individual annotators have highly variable and changing styles, no common palette of symbols or collection of text boxes will be able to support the rich vocabulary of *ad hoc* arrows, asterisks and brackets, and, more importantly, the fluid and opportunistic transition between them [43]. This fluidity should apply not only within a single document but also between documents [74].

- Many readers use informal coding systems in which colours of ink or symbols carry specific meanings. Coding systems should remain informal, only going as far as is useful to each author, and their presence should not be imposed on those readers who do not use them. [43]

- Support for the distinction and curation of public and private annotations is important. Several studies [43], [89], [46] have noted the value of a reader's annotations to others, as prompts, feedback or in collaboration. A significant issue here is the addition of time and resources required to convert between private and public annotations [46].

- Seamless integration of annotation functionality within the reading experience is the most important feature to support according to Marshall [43]. Mode switching between reading and writing costs time and attention [78], and readers do not do it willingly [76] or well [59]. Instead, they annotate with whatever tools are available [43]. Maintaining engagement with the text while annotating is important and difficult to preserve.

Digital annotations need to conform to a long list of must-have features to match the ubiquitous, in situ, and seamless annotation techniques seen on paper. In addition, digital annotation may also greatly improve upon the range of uses and power of paper annotations, and the digital realm has several key advantages over paper. E-paper or OLED tablets are rigid and marking them up is possible even in the absence of a desk or a flat surface [15]. While on paper annotations are scattered throughout the document by virtue of being inextricably linked to their anchors, in the digital realm annotations may also be aggregated and displayed elsewhere, for instance on a summary page where highlighted excerpts and reader-generated notes can become a standalone study aid [75]. Additionally, digital annotations may take up an arbitrarily large amount of space instead of being restricted by the size of a printed document's margins, they can keep track of their authors, timestamps and expiry periods as required, and they may be searched, filtered, synchronized across devices or shared.

3.2 Navigation and Orientation

In text documents, the vast majority of which span more than a single page or view, **navigation** is the process of bringing into view and guiding one's attention to the parts of a text relevant to the user's current reading goals. Navigation is a vital activity that all readers perform, and it encompasses a wide range of tasks and actions, depending on the context and the type of reading performed. We distinguish navigation as a reading-supporting activity from reading itself. For instance, scanning a page or view in search of a specific word does not involve navigation until the page needs to be changed. **Orientation**, by contrast, is the awareness of the structure and nature of the text and the reader's position within it. If orientation is to be thought of as awareness of location, so navigation is analogous to displacement or changing that location.

The naive view of reading regards navigation as the occasional turning of a single page as the reader reaches the end of the current one. However, in a study on magazine reading Marshall and Bly [45] identify a wide range of navigation activities and techniques that reveal a far more nuanced picture: participants unselfconsciously and quickly flipped through pages, glanced back to reread or forward to see how much was left in an article and skipped ahead to articles of interest identified in the table of contents. Navigation is an essential companion activity, especially to responsive reading, which is far less linear than casual or receptive reading.

We identify two different kinds of navigation: *intra-document* and *inter-document*. Intra-document navigation takes place within the confines of a single text document. In contrast, inter-document navigation refers to the management *of* documents while reading, determining which one to attend to next and the awareness of each document's purpose and contribution to the goals of the reader. This boundary is not rigid, given that multiple chapters in a book may be considered different documents even if they are part of the same printed volume. Although not explicitly discussed, Marshall and Bly's [45] magazine reading participants employed elements of intra-document (managing attention within articles) and inter-document (switching or navigating through unrelated units of text) navigation.

The main reason for this differentiation is to elicit awareness of the fact that it is inter-document, not intra-document navigation, that raises more issues in digital responsive reading. In Tashman's [74] study on responsive reading, about 16% of issues cited by readers of digital text had to do with the poverty of representing multiple documents in a meaningful way and enabling the effortless switching between them. In contrast, only 5% of issues reported were with intra-document navigation. Additionally, most knowledge work tasks involve synthesizing information from multiple sources to arrive at a coherent position. This means that we should consider inter-document navigation an essential component of the problem and regard **multi-document** reading scenarios as the *default* for responsive reading [74].

In a freely configurable multi-document workspace, inter-document navigation takes on another aspect of work organization. Not only is the reader in charge of which portions of text are visible on each document, but also where each document sits on the desk, which documents are more easily accessible, etc. [45] This aspect of navigation, known as **workspace layout**, is a central workspace-level requirement of responsive reading.

3.2.1 Components and Attributes

We identify the following attributes of a navigation action:

Trigger

The *trigger* of a navigation action is the reason why performing it becomes imminent or desirable to the reader at a particular moment. The most common navigation trigger is linear: reaching the end of a page while reading necessitates the turning of the page. Other, non-linear triggers, result from the desire to read a different part of a document, or a different document altogether and occur with varying frequencies depending on the type and the goals of the reading.

Interaction

The *interaction* involved with a navigation action is the manipulation of the text undertaken to accomplish the action. On paper, this could involve turning to the next page, rifling through a book's pages or folding a page out of the way to expose another one [45]. Digitally, these interactions vary wildly in modality, mode and maintenance of flow and the extent to which they interrupt the reading task.

Target

The *target* of a navigation action is the final view the reader aims to achieve through the interaction. This target may be explicitly known (e.g. "the passage two pages earlier where that character was mentioned"), it may be explicitly identifiable but not previously seen (e.g. "see section 1.4"), or it may be neither (e.g. “they must detail their methods somewhere”). Finally, the reader’s purpose may not actually be known and a reader may initiate a navigation action in order to serendipitously come across something of interest [45].

Breadcrumbing and Duration

Navigation actions differ in their *durations*, or, more accurately, in the amount of time that passes on a single view between discrete navigation actions. In receptive reading, simple page turning is the most common navigation action, and turning pages is expected to be interspersed with long periods of reading characterized by navigational inaction. In contrast, search reading or summarization may require frequent navigation actions with the resulting views remaining in focus for very short durations. Flipping back and forth between two conflicting passages is an example of such short-duration, almost oscillating, navigation actions that effectively cancel each other out in pairs. If a short-duration action is initiated, the reader often has the means to quickly reverse it (e.g., using a finger as a bookmark before flipping forward). This form of unconscious breadcrumbing and reversability of navigation actions should be supported digitally as well.

3.2.2 Navigation-Aiding Features in Text

Navigating a printed document has not changed drastically since the bound codex (the precursor to the present-day book consisting of bound, turnable pages) gained prominence in the 3rd century C.E., but new features to facilitate navigation have been gradually incorporated and standardized.

The navigation-aiding features in a printed text (not all of which may be included) are:

- **Page numbers** are unique, sequential and unrelated to the content of the text. They are useful in locating specific pages and in aiding the reader’s orientation (See Section 3.3), but they are not linked to the semantic structure of the text. Page numbers are fragile references that will be different in different editions of the same book, and repaginating a digital document would reset any location cues they may have provided.

- **Sub-document headings** like chapters, sections or verses may be named or numbered, breaking up the text into smaller, typically hierarchical, identifiable chunks. They are directly linked to the meaning and structure of the text, which persist regardless of the format of the document. The Holy Bible is an example of a deep hierarchical structure, and it is divided into groups, then books, then chapters, then verses. The former two are named while the latter two are numbered.

- A **table of contents** combines the previous two features to provide an at-a-glance outline of the semantic structure of the text with clear page number correspondences for each section.

- **Indices** match concepts, names or words to places where they appear in the text (typically by page number). An index provides a one-to-many mapping between a concept and its occurrences in the text. On paper, a downside to indices is that they must be manually created, adding an interpretation step to what the index-worthy concepts or names are for a particular text, meaning that a manual index is by nature interpreted, curated and incomplete: it cannot help locate every possible concept in a text.

- **Physical features** such as bookmarks or recessed dictionary tabs may be created at printing time or placed by the reader to locate a particular page. Their meaning can be rigidly determined as is the case with dictionary tabs, or it may be completely arbitrary.

Electronic navigation can incorporate most of these navigation features, while empowering new ones with a complete awareness of the text and library-wide accessibility unencumbered by physical limitations.

- **Search** removes the need for manual indexing by keeping an idea of the location of every occurrence of every word or part of a word in the text. This means that as long as the reader can provide a search string, she can quickly find and navigate to any of its locations.

- **Hyperlinks** offer the ability to select a section of text with a previously assigned redirect and immediately be taken to a correspondingly linked portion of a document. Links were originally used on paper in the context of interactive fiction, and with in-text redirects such as "see Section 1.4.", but in the digital realm they are far more advanced: they can be followed immediately and effortlessly without the need to keep track of unrelated location information such as page or section numbers. Hyperlinks and hypertext are discussed in greater detail in Section 3.4.

3.2.3 Immediacy, Unconscious Use and Serendipity

In several comparison studies of paper and digital reading [45] [59], the qualities of navigation most cited as superior in the paper condition were its ease of use, integration into a reading flow and immediacy of feedback. Navigation on paper is rapid, unselfconscious and well-understood by readers. Readers manipulate pages as they require, in a continuous manner, interleaving navigation activities with reading, annotation, etc. In addition, paper navigation is bimanual, meaning the primary hand need not let go of writing implements to navigate, and it is tactile, enabling navigation activities to be performed without taking visual attention away from the task of reading.

Many attempts at electronic navigation have lost some or all of these qualities, requiring dominant-hand navigation input, locating a button in a visual interface, or presenting a limited configuration of views. For instance, on paper a page can be partly bent out of the way to briefly expose the page below without disappearing from view. In contrast, most digital reading appliances implement page navigation in a discrete manner. Only recently, on the recommendation of Marshall [45] has visual page turning feedback in appliances like the Apple iPad been made continuous in order to keep two adjacent pages in view. Users reading webpages suffer even more, with the continuous vertical scroll metaphor making it extremely difficult to see two passages of the same document side by side.

Finally, e-book appliances with e-paper displays introduce a detrimental latency to the feedback of navigation actions, which requires them to be implemented discretely. E-paper screens refresh slowly (1000-2000 ms), a whole screen at a time, and e-paper documents cannot be riffled through as quickly as bound books can. This interferes with serendipitous navigation (flipping through a magazine until something interesting catches one's eye), rapid scanning and search, and using images as anchors to help find passages of interest.

It should be noted that many of the activities described above are possible on most digital reading appliances. Given a large enough screen, a PC can display the same document in two windows side by side and enable independent navigation on each copy. However, this is highly dependent on screen space [59], and requires a mode switch, at least in the reader's mind. The navigation action has to be formulated in its entirety before it can be executed, increasing the cognitive load on the reader and breaking the reading flow. Even so, there are examples of versatile reading systems such as Chen's [16] "United Slates" system that mitigate some of these issues by making multi-view navigation more intuitive.

3.2.4 Orientation

Orientation is closely related to navigation since every navigation action needs to synergistically update the reader's awareness of position. The views and triggers of navigation actions depend on the reader's orientation within the text. Orientation, in turn, relies on an internal representation of the form and structure of the text: how much is left to read, of what important sections does the text consist, etc.

Some participants in Marshall and Bly's magazine study [45] used the table of contents in service of orientation rather than navigation, building an internal awareness of the document but not necessarily flipping forward to a particular page number.

There are many orientation cues in printed materials. The simplest of them is a tactile one: how far into a book the reader is can be immediately gauged by the relative size of the portions of the book in her right and left hands. The pages themselves also offer orientation cues and can affect the navigation action: flipping three pages forward feels different than flipping three hundred pages forward. Navigation actions may also be undertaken for the purposes of orientation, with a reader flipping through the remainder of a document quickly before returning to the original reading spot.

Many of these cues and contexts depend on the notion of a page, which is far more rigid in printed text. In the digital realm, the page concept is very broad. A webpage may be arbitrarily long, while an e-book page on a smartphone screen may only span a few sentences, which contributes to confusion and undermines orientation. Participants reading on a screen where a page may be arbitrarily large have been shown to quickly become frustrated and feel disoriented [49]. Orientation is further hampered by the fact that in digital documents, the notion of page may change in the middle of a reading: changing the font size repaginates the document and turns a 300-page document into a 500-page one without changing its content, all the while obliterating the layout of the previous set of pages and destroying any orientation cues the reader took from the relative position of graphics or passages to the boundaries of the page. The issues of layout and cognitive mapping are discussed in the following section.

3.3 Layout, Cognitive Mapping and Workspaces

The **layout** of a document concerns the way its visible components are arranged, and includes considerations of pagination (real or virtual) and the format and placement of elements such as headings, graphics and text blocks. A consistent layout improves orientation and facilitates activities such as **cognitive mapping** and recall. The layout of a **workspace**, by contrast, concerns how different documents are arranged and can aid in high-level activities such as document search, triage, and reading from multiple documents. These companion activities all rely heavily on the cognitive structure and properties of visuospatial working memory.

3.3.1 Visuospatial Memory

Decoding spoken language relies on a working memory structure called the auditory loop. This is the mind's inner monologue, the buffer one uses to repeat a friend's phone number until it is memorized. The auditory loop is linear, susceptible to interruption, and low-bandwidth. The task of decoding written language typically also engages this structure.

In contrast, reasoning about spatial data occurs in the visuospatial sketchpad, a staging area that is tuned to manipulate two-dimensional information and efficiently encode images to long-term memory. It is highly developed due to the requirements of mental navigation in early plains-dwelling humans and it greatly predates language. In fact, visuospatial memory is more efficient than auditory memory due to the mind's ability to easily encode mental pictures along two dimensions. This facility has been recognized for millennia and has been consciously adapted in the service of non-spatial memorization. Ancient orators used a mnemonic technique known as the "method of loci" [18] to remember speeches by placing representations of concepts in an imaginary two-dimensional environment. A version of this technique is used by competitive memorizers to recall vast amounts of non-spatial information such as the sequence of cards in a randomly shuffled deck.

While the naive task of reading (i.e., decoding written language) does not employ the visuospatial sketchpad, there is compelling evidence that the properties of visuospatial memory are constantly leveraged in the broader tasks of responsive reading. This occurs both at the document level, where a passage might be found by its relative location to page boundaries, and at the workspace level, where spatially arranging documents helps with triage and organization.

3.3.2 Cognitive Maps

In a seminal exploration of visuospatial memory Edward Tolman coined the phrase "cognitive mapping" [79]: a method for constructing and accumulating spatial knowledge so that it can be retrieved through mental visualization. Initially studied in rats, cognitive mapping originally referred specifically to storing information about "everyday spatial environment[s]" [34]. However, the term has expanded to refer to any method of storing information that has a spatial component in the mind, regardless of whether the original information is spatial in nature.

Waller [86] synthesizes that responsive reading helps readers build a "cognitive map" of the text they are reading as a physical object, using features of the text such as figures and headings as landmarks. This is referred to as **spatial mapping** of text. This mapping allows the reader to visualize the layout of the text, to correctly place headings and figures within it and even to remember the approximate location on the page of a particular sentence or definition. It uses page boundaries, as well as illustration and other non-textual elements as salient landmarks for the content around them. In a study of this phenomenon, Rothkopf [65] found that participants could remember the approximate part of the page on which a passage was to be found significantly better than chance. This is a shallow, subconscious mapping that can be created regardless of the content and semantic structure of the text. It is useful when looking for a particular passage to consult or as an unlikely ally when trying to remember content absent the source material, but it is unreliable and fragile, disintegrating easily with a change of layout or font size.

In contrast to spatial mapping, which is concerned with the visual elements in a text, there is some evidence that readers also engage in **conceptual mapping**: identifying the key concepts in a document as per van Doren and Adler's [2] third reading level and arranging them in two-dimensional space to capture the relationships between them. Since the spatial characteristics of written text hold no salient meaning by themselves, the particulars of the text's layout may not be as important as the way the concepts and ideas described therein are interlinked. Conceptual mapping is the process of forming a sense-making two-dimensional representation of the relationships between the ideas in a text. It also utilizes the mnemonic properties of visuospatial memory, but unlike spatial mapping, conceptual mapping persists with different editions and layouts of the document, because it is a mapping of the content, not the layout, of the text.

Since conceptual mapping is a content creation activity, it has specific interaction requirements. It is discussed in greater detail in Section 3.4., where a learned technique called mind-mapping is described.

3.3.3 Consistency

Much like dentists have consistent tool placements so that they can reach for each instrument without having to identify it consciously, the fact that the layout of a paper book is fixed the moment it is printed is useful to readers. A reader relies on the constancy and consistency of a physical book, both for the aforementioned companion activities, and in a more metaphysical sense. A printed book will never change unpredictably due to updates or repagination. Absent physical damage, the book will retain the text in the exact place it has always been.

The lack of consistency is a significant obstacle to forming spatial maps of digital text. Even if the text is divided into pages (which, in the absence of graphics, are the only sources of visuospatial cues with respect to text), this division may change multiple times due to changes in font size or typeface. Worse yet, if a document is presented in scroll mode (e.g., as a webpage), it loses even more of the possible cues to spatial mapping [59].

Consistency is important because visuospatial information is encoded in memory with the presumption that it will not change without conscious effort. The original purpose of the visuospatial apparatus, namely to enable our ancestors to navigate mostly flat terrain, presumed consistency because the spatial arrangement of landmarks in the world was constant. Even if a landmark looked different or changed over time, its location would remain the same. Any alteration to the source renders the memory representation useless. We see visuospatial consistency in use in desktop and mobile interaction conventions (e.g., the location of the Start button on a desktop or the Home button on a smartphone), and by individual users both digitally (e.g., PC desktop arrangements) and physically (e.g., workspace arrangements).

3.3.4 Workspace Layout

A behaviour noted by O'Hara and Sellen [59] in their comparison of online and paper reading was that users often encoded meaningful information in the spatial arrangement of documents on physical surfaces, and were unhappy with the lack of support for such behaviour on their desktop computers. Switching between reading and writing and navigating through multiple documents were far easier on paper on account of this affordance. Rather than cycling through consecutive windows, readers at a physical desk could easily recall the position of a document and reach for it. Additionally freeform workspace layout allows readers to designate arbitrary workflows and locations for documents, and perform efficient document triage and sorting. In fact, the ability to freely arrange documents in one's workspace is arguably more important than the incidental memory for text demonstrated in spatial mapping.

Due to the limitations of screen size and orientation (which is predominantly widescreen), most desktop and laptop computers have limited support for freeform spatial arrangements of application windows, as well as a singular notion of focus. This can be mitigated in one of two ways: by enlarging the digital workspace or by rearranging devices containing different digital views in a physical workspace.

Since the digital workspace is limited mostly by screen size, the enlargement approach alleviates the problem by increasing screen size. In PC's, this is achieved by utilizing multiple monitors to create a larger working desktop. More specialized devices such as tabletop computers are predicated on the idea of a large, potentially sharable and collaborative digital workspace that more closely resembles a physical desk. In either case, the reader is able to spatially arrange digital documents and views on a large digital display. In addition, desktop prototypes such as BumpTop [3] enabled the freeform spatial organization of arbitrary chunks of digital information such as photos, application shortcuts, files and folders.

The combined physical-digital approach leverages the kinesthetic cues of multiple physical objects by retaining the single-view metaphor for each digital device while having multiple devices arranged spatially on a physical desktop. While the devices need not be linked, work by Chen [16] and Hamilton [28] have made these devices aware of each other and have explored the challenges of device navigation. Chen's [16] work was specifically geared towards responsive reading. There are additional benefits to the blending approach as there is evidence [73] that kinesthetic cues facilitate the encoding of visuospatial information.

Multi-Session Preservation

As part of the consistency requirement for visuospatial memory, the layout of reading workspaces should be preserved between sessions. This is easily achievable on a physical desk, where documents will remain as the reader left them without the need to save or manage them. However, according to a study on desk organization by Malone [41], arranging one's desk in a meaningful way has a high cognitive cost, and it serves the dual purpose of helping locate documents and reminding workers of action items that need to be performed. If knowledge workers put the same amount of effort into organizing their virtual workspaces, they should remain consistent between sessions.

3.4 Non-Linearity and Hypertext

The reading of text at the decoding level is inherently linear. Symbols are decoded one after the other in an order established by the directionality of the writing system, be it left-to-right, right-to-left, or top-to-bottom. The linearity of text may be reasonably extended to literary narrative, which is expected to be read from beginning to end. However, from our discussion of reading comprehension we know that mental representations of text are anything but linear. Concepts in the mental landscape may be activated in any order, triggered not only by the current phrase being read, but also by cohort activation of similar concepts, by prior knowledge, or by deliberate external research.

At best, even if the text is being read in a strictly unidirectional way, it requires concept activations from a non-linear network, with each phrase's concepts extending links to previously read passages, to other texts, or to the reader's general knowledge. In fact, text ceases to be linear the moment it is represented in the mind.

Nowhere is this more evident than in a scholarly article. Even the document itself, still subject to linear decoding at the lowest level, deliberately shifts the reader's attention from the main text to footnotes or to the references section and back, requires parallel understanding of a figure and the accompanying text, and is subject to responsive reading strategies such as skimming, scanning or glancing back, all of which defy the plain linear reading expected of short, simple texts. The path taken through such a document is never linear, and may be different for different readers.

We have already established the mental representation of text as an interconnected graph of concepts. Rather than conceiving of a document as linear, we must understand that its text also forms an ad hoc graph of textual nodes that a reader may enter, traverse, and exit in a potentially infinite number of reasonable ways [26]. In fact, we will apply the graph metaphor at two additional levels: within-document and between-document.

The **within-document graph** consists of nodes of text, each containing the result of a single linear symbol decoding episode. When the reader stops decoding and changes location, she forms an ad hoc link to another textual node in the graph. Some of these changes are triggered metacognitively by the reader (e.g., "I should re-read the thesis paragraph."), while some are deliberately set by the author (e.g., "See Section 11.3"). Eventually, through the graph of near-infinite possible nodes and jumps, the reader traverses a single linear path, the path of what he or she actually read in chronological order.

Since a reader can almost never acquire all relevant knowledge or understanding from a single document, the path one takes through a sequence of documents also matters. The **between-document graph** emerges when a reader is directed from the current document to others. In a literary trilogy, for example, the first document implicitly points the reader to the second one, which points to the third. Scholarly articles end with lists of other documents relevant to the reader. Granted, a deliberate list does not exhaust all possibilities, as the reader may choose to open any document available to him if he believes it will be beneficial to the task. In this ad hoc graph, each document may link the reader to a potentially infinite number of related documents, implicitly or explicitly, and the reader traverses a linear path through documents in a fashion similar to the one described in the within-document graph.

In fact, it may turn out that a distinction so recently and carefully made is in fact meaningless. Given that a document may be arbitrarily long, what if we called each of the blocks of linearly decoded text from the within-document graph a document in its own right? Then a reader's linear graph would consist of consecutive blocks of text that link to each other in some way, and the within-document and between-document graphs would be unified into a single graph of text blocks and links between them.

Since a reader's path through a text will likely not be linear, care must be taken to provide the tools necessary for his or her continued orientation, support for navigation and, most importantly, support for quickly and easily returning to the source text after a jump, if required. Paper documents do not excel at this, as even explicit between-document links place the onus of finding a document, retrieving it, and keeping one's place in the original text, squarely on the reader. However, paper makes it possible to see two parts of a document at once through simply bending a page or moving it aside. Multiple documents may be spread over a reading area, making their simultaneous or alternating viewing simple and effortless. In a world of digital technologies that overwhelmingly mimic traditional reading by displaying full-screen single pages, navigation within and between texts is slower, requires deliberate cognitive effort, and may in fact become less frequent over time, resulting in the suppression of non-linear navigation.

However, despite digital technology failing to support non-linearity on scanned paper documents, it is vastly superior in another aspect of digital reading: hypertext.

3.4.1 Hypertext

Although the potential to access a document's text blocks out of sequence has always existed in reading, the advent of digital technology has made it explicitly customizable through the concept of hypertext. Hypertext, a term coined by Nelson [57] meaning "more than text", enables links to be embedded in a document, deliberately connecting to other documents. More precisely, as defined by Crane and Mylonas, ""hypertext" refers to the electronic linking of blocks of text" [20]. On an interactive digital display, accessing the link results in the target block appearing on the display, often in place of the source block. Hypertext enables the creation of "conceptual and literal links among disparate sections of a given text or among completely separate texts" [24]. With hypertext, both the notion of a single document and the distinction between documents are altered in the extreme.

Hypertext, as understood above, has led to the vast changes in reading patterns and behaviour. The information exchange infrastructure known as the Internet has enabled the creation and utilization of the World Wide Web, a system of interlinked hypertext documents [12]. Despite the WWW expanding to encompass and support linking to other media (images, video, sound, programs), at its core it consists of a vast, deliberately linked network of document blocks. Since its inception, the WWW has dispensed with several forms of printed material (e.g., dictionaries, thesauri, encyclopedias) and is a favourite means of distribution of previously non-hyperlinked content such as news and magazines. The WWW has also drastically affected reading patterns: we read more summaries and shorter texts, and more text of uncurated and unverified origin and correctness [39]. More importantly, the path we take through the text includes following more links between shorter text blocks, making the issues of navigation, orientation and breadcrumbing (i.e., the ability to retrace one's steps) more important than ever.

Writing Hypertext

When a traditional, non-hyperlinked document is written, the author has the opportuity to introduce, develop and conclude a linear argument from beginning to end. The hierarchical form of organization in which the introduction-body-conclusion cycle may occur on several levels of detail. The overwhelming majority of legacy content, and much of what is being written today, conforms to this linear structure. With it comes the expectation that the reader will more or less traverse the document in such a way as to follow the author's intended path.

Writers of hypertext documents may wish to retain this hierarchical structure, but they are also able to weave other mechanisms of organization threads through the content by using different types of links. The onus of finding a single path through hypertext rests overwhelmingly on the reader, meaning that the writer is less responsible for presenting a linearly coherent text and more responsible for laying out a hyperlinked landscape in which many such paths are possible and valuable to the reader. This introduces a new dimension of opportunity and challenge for the writer, as it is impossible to generate the full complement of links between blocks, and deciding which ones are included constitutes another level of authorship that may affect the message of the text.

The reader, by virtue of following links he or she deems relevant, reads a body of "self-authoring text" [7] of his or her own creation and curation. Although the writer cannot control which link the reader will follow at any given time, he or she can curate which links are in fact available to follow, and the exact span of text to which they link. According to Fitzgibbons [24], this makes hypertext especially suitable to scholarly or literary communication as it enables the deliberate organization of a document on many different levels alongside the original linear structure.

Reading Hypertext

Users of the WWW are logging countless hours of experience reading hypertext. Encyclopedia link hopping, news aggregators and social media all rely on the concept. However, returning to the focus of this document, how much of hypertext reading is or can be responsive? In other words, how does hypertext reading support the acquisition and consolidation of new knowledge?

There has been much speculation and work on the role of hypertext in education, even though so far no definitive process or strategy has emerged. The crucial seed of potential in hypertext lies in the following, as described by McKnight [52]: education is a process not of memorizing and regurgitating linear information, but of forming a usable network of knowledge and thinking skills that may be applied to further a particular field of study or enterprise [10]. As such, it is paradoxical that the traditional methods of education rely on linear content (textbooks, lecture slides, course prerequisites) to foster non-linear skills. Hypertext has at least the potential to empower the learner, as a reader, to follow a non-linear path through material in a particular domain.

Protopsaltis and Bouki [62] developed a model describing the process of reading a hypertext document with a particular focus on comprehension. They contend that readers interact with hypertext both in linear and in non-linear ways, and that the driving force behind hypertext reading is, as with other kinds of reading, the satisfaction of a reading comprehension goal. This employs the metacognitive apparatus described in Section 1.1.4. and consists of repeating cycles of scanning the links available, deciding on a strategy, building and incorporating knowledge, and evaluating progress.

The Bibliographic Universe

The Memex, as described by Bush in 1945 [14], is a mechanical extension of human memory in which all documents a person deems important are stored and interlinked so that a single concept's occurrence may be traced through multiple documents, transcending their boundaries. Bush envisioned the Memex as a highly individualized and personalized structure, but there is tremendous value in its more general counterpart. What if this structure contained *all* documents in the known universe and the links between them according to any concept or search requirement? This would be the Holy Grail for library science, a system that does not pigeonhole a document inside one hierarchical cell of human knowledge, but positions it precisely in the hyperspace of what Fitzgibbons calls "the bibliographic universe" [24], linked to, and accessible from, all other documents that it relates to. The impulse to collect all the world's written knowledge in a universal library dates back to Alexandria, but it is only with the advent of hypertext and digital storage that such universal catalogues have been made possible.

Navigating Hypertext

The difficulties and requirements of multi-document navigation have been outlined in Section 3.2. These challenges are further exacerbated by the shorter length of hypertext documents and the general trend of single-document views that pervade modern web browsers. Given more documents and more hops, without a tremendous amount of additional orientation efforts, the ability to revisit a previously read passage or to retrace one's steps generally is extremely limited [24]. A significant downside to hypertext reading is that each "self-authoring text" exists only as long as each hop is remembered and traceable, either in the reader's memory, or through the use of breadcrumbing and link histories. A single-page overview of a multi-dimensional linked document is also much more difficult, as is an initial skimming. Searching for information by following links in a badly organized or badly presented hypertext document can turn into a combinatorial nightmare. The collision of non-linear documents and linear traversals may leave the reader completely lost and disoriented in the absence of sufficient navigational support. While even a cursory exploration of hypertext navigation is beyond the scope of this document, we recognize that as responsive hypertext reading becomes more prevalent, small innovations in this area may have a tremendous impact on the suitability of hypertext environments to responsive reading, and to utilizing the bibliographic universe to its full potential for knowledge access.

3.5 Other Related Activities

In this section we briefly mention several companion activities to reading that we believe are important, but exist slightly outside the main thrust of this document. The main reason for this distinction is that unlike annotation and navigation, which are companion activities intrinsically tied to the text and the technologies used to access it, the activities in this section, while performed alongside reading, need not use the same technology as the one that displays the text. Notetaking, for instance, may be performed as easily alongside a paper document as it can alongside an e-reader digital version, on a legal pad, laptop or tablet of the reader's choice. This decoupling allows us to label these activities as less urgently in need of innovation and intervention.

3.5.1 Notetaking, Clipping and Summarization

Clipping refers to the ability to lift entire quotations from the source material and insert them into related written artifacts. This is useful in the context of **summarization**, which entails creating a written document to be consulted instead of the source text in the future. Different readers will find different kinds of summaries useful, and will therefore generate different kinds of summaries depending on the situation and their goals with respect to the text. However, all summaries require the distillation of the contents of the original text. As such, summarization is a learned skill that entails competent use of strategies and metacognitive outcome monitoring. While studying the process of developing these strategies and employing macrorules in summarization, Brown and Day [13] noted the better trained learners had access to more sophisticated summarization and composition strategies and used them to produce more accurate and useful summaries.

The summary of a document, when compiled correctly, can in fact replace the source text in the reader's future research activities. If the reader has to engage with a large number of documents, reading their summaries first would be valuable in reading triage. This is one of the roles of abstracts in academic publications. A considerable body of work in natural language processing has focused on automatic text summarization (see [35] and [42] for examples), hoping to relieve the reader from this cognitively demanding task.

However, more than the end product, the process of summarization itself is valuable to responsive readers. The additional cognitive load of performing summarization is offset by the fact that the task forces the reader to make judgments about what is important enough to include in a summary, in effect reasoning about the conceptual structure, overall thesis and importance in the domain of the document.

Duke and Pearson synthesize that "[t]eaching students to summarize what they read is another way to improve their overall comprehension of text" [21].

Notetaking is a content creation activity similar to summarization, but it typically focuses on information that is not available in the source document, including the reader's own interpretation or that of a third party such as an instructor. Like annotation, notetaking occurs while a text is being read, but unlike annotation, which is marking up the source document itself, notetaking is typically performed in a different view. Reading platforms need not support notetaking as long as they can integrate and synchronize with the reader's preferred method of creating external content. This becomes difficult in a mixed-media (digital/paper) situation, despite efforts by various companies to achieve seamless integration between the two using digitizer pens or specially printed scannable markings on paper. With the creation of additional content through notetaking and summarization, the ability to manage these documents (if created on the same device) and to consolidate them (if created on different devices) becomes crucially important.

Mind Mapping

In Section 2.3, we described the internal representation of concepts in a reader's mind as a graph of nodes. A popular visualization strategy called "mind mapping" is a deliberate, often explicitly taught summarization method that trains users to represent the relationships between concepts spatially on paper as they are reading a document or listening to a lecture. In mind mapping, rather than taking notes chronologically, readers extract and arrange key concepts in a spatial diagram, with the main premise or concept in the centre. Related concepts are linked with lines of varying colours and thicknesses, and the result is a single-page two-dimensional summary of the relations between ideas in the document. Due to the fact that concepts do not necessarily form a spatial arrangement in the source text or in the mind, mind mapping must be done consciously and strategically by the reader, and requires precise, freeform drawing tools. Mind mapping has comprehension and summarization benefits and has been adopted by scholars and professionals in a variety of fields [53].

3.5.2 Collaboration and Sharing

Collaboration while reading is an important activity to support, and a difficult one to describe. Collaboration may be collocated or not, synchronous or not, and it may involve multiple actors performing any of the other activities discussed in this section. That said, there is a strong indication that collaborative reading is a fundamentally different kind of activity from individual reading, made more complex by issues of communication and negotiation, privacy and authorship of the generated written artifacts and the cultivation of a public persona who might behave very differently when reading in a group. A full overview of the collaborative reading literature is beyond the scope of this document, as it will focus on multi-document, multi-technology, but single-reader scenarios and interventions.

However, there is an important public aspect to responsive reading that concerns the management and sharing not only of the source document, but also of the additional content the reader generated. As discussed in Section 3.1., not all content a reader generates as part of the reading task is intended for, or fit for, public consumption. Different levels of privacy may be important to the user. [46]

Chapter 4

Researching the Practices of Academic Reading

Academic reading is an important subset of responsive reading. It is the type most commonly employed by students, scholars and academics when they are learning not only domain-specific information, but also meta-cognitive strategies about learning itself.

There are some special considerations to take into account when conducting research on academic reading. Academic readers are a very diverse group in their relationship to technology. Rather than dividing them along generational lines or labelling them digital "natives" or "immigrants", we propose that academic readers engage with technology in a variety of ways and with a variety of motivations that depend on the specifics of their tasks. Their technology usage profiles are the result of many years of successful responsive reading that has been tuned to each individual's needs and preferences.

In contrast, the approach to researching new technologies for academic reading has been largely impositional: the typical academic reading study, especially at the institutional level, designs and deploys a slightly modified mainstream device to a cross-section of students and collects data and feedback that is used to determine the suitability of the device as an alternative to paper for all students. This may be the only way that such large institutions see the advancement of digital academic reading, but it fails to account for the aforementioned diversity in readers. We believe that each individual academic reader is already using his or her own unique set of habits, workflows and strategies to great success, and that the collection of these individual preferences and behaviours constitute an emergent set of reading competencies from which much can be learned.

Another shortcoming of much of the research and development to date is to attempt to replace or dismiss not only paper, but also every other kind of digital device as the "best", the "only", or the "most suitable" for academic reading. The reality is that with the increased proliferation of the aforementioned variety of digital devices, we must consider the default technological scenario to be that of integrative use: how multiple heterogeneous devices could be managed and put in the service of the same overarching goal or workflow in the context of academic reading.

4.1 Academic Readers

Academic readers are knowledge workers whose primary professional reading tasks involve academic materials. Even in this small subset of knowledge workers we can identify at least three distinct groups, which have very different reading profiles. We deliberately exclude primary and secondary school students from this analysis, for two reasons. One, students below the age of eighteen do not have full decision-making freedom with respect to their device ownership and usage. Their workflows are often highly prescribed and constrained. Two, no assumptions can be made about the reading abilities of this demographic as they are, for the most part, obligated to attend school and acquiring reading skills at an inconsistent and variable pace. We shall consider academic readers to be only those who have chosen post-secondary education as adults.

4.1.1 Demographics

We shall consider three groups of academic readers:

Students are undergraduates whose academic reading is almost always conducted to satisfy some external evaluation criterion. Their reading materials are typically selected for them by their instructors. In addition, the coupling of reading and feedback is quite close, meaning that reading each document is soon followed by a predictable evaluation of the concepts learned in the document. Undergraduates are primarily 18-25 years old and represent the most digitally entangled group of the three: generally, their school-related workflows have utilized digital technology for a longer portion of their academic development.

Postgraduates are students pursuing master's degrees or doctorates. While still learning and improving their academic skills, their reading is increasingly characterized by content curation and triage. They have a greater freedom in selecting relevant documents, meaning they are responsible for deciding not only how, but what to read. An additional constraint is that evaluations of their reading strategies and abilities may be fewer and less directly tied to individual documents. Postgraduates may also take on some mentorship and evaluation roles, reading others' documents critically for the purposes of grading or peer review and having the power to recommend changes. Postgraduates vary widely in age, but few are likely to be younger than twenty-two.

Scholars are readers who have chosen academia as a profession: postdoctoral fellows, researchers and professors. They are no longer considered to be in the training phase of their careers, although they have signed up for lifelong learning and improvement in their chosen field. The evaluation of their reading tasks is the most nebulous and long-term out of all three groups, and they are active mentors of students and postgraduates. Generally, readers do not become academics much before the age of thirty, meaning that few academics would have used Wikipedia in high school.

4.1.2 Relationship to Technology

What relationship do academic readers have with digital technology? If the groups above are taken naively, stereotype would suggest that academics are generally out of date and resistant to technology while students feel perfectly at home in the digital realm. This contrasting of students and academics as somehow belonging to categorically different generations with respect to digital technology gave rise to the "Digital Natives/Immigrants" hypothesis. First proposed by Prensky in 2001 [61], this hypothesis states that the then-current generation of students had grown up with digital technology and was therefore going to be irreconcilable with the generation of their instructors, who were not native users of digital technology, but rather digital immigrants who could never fully adapt to the new landscape. It also implicated traditional educational paradigms as being outdated and unsuitable to instruct digital natives.

This distinction cuts broadly across generational lines, and has ceased to be useful in the subsequent decade due to severe criticism and marked changes in the digital landscape. One could argue that the digital natives of 2001 are now on the verge of becoming academics themselves, ready to embrace the digital revolution in education alongside their digital native students of the following generation. Alternatively, one could point out that cohort-wide generalizations exclude so many differing cases as to hold no descriptive or discerning power. Not every student born after 1985 has had the same exposure to digital technology, due to anything from economic factors to personal preference. Even those who have may not be comfortable with, or necessarily prefer, any arbitrary digital education tool over its legacy counterpart. Finally, there is no evidence that the supposedly untrainable digital immigrants could not keep pace with technology and apply it to their teaching. The natives/immigrants generational divide is no longer applicable.

The factors to consider in the relationship of academic readers to digital technology are, fortunately and unfortunately, more nuanced and more complex. Another topology, proposed by White and Le Cornu [87] differentiates between users on the basis of their online engagement for specific tasks and defines a continuous spectrum between digital visitors and digital residents. White and Le Cornu's digital visitors see the Web as "an untidy toolshed" [87]: a repository of tools that they may need to accomplish specific tasks. While visitors will enter the shed, select the appropriate tool, and use it, they will not linger online, because the Web is not a hospitable place in their minds. Digital residents, in contrast, "see the Web as a...park or building...in which there are clusters of friends and colleagues" [87]. A significant part of a digital resident's life is actually lived online. Friendships are formed, experiences are shared and collaboration is accomplished. This characterization offers several advantages over Prensky's. Rather than consisting of binary categories, the visitors/residents axis is a continuum. In addition, an individual user may lie in different places on the axis for different task or behaviours (e.g., a user who communicates comfortably online but reads overwhelmingly on paper). It makes distinctions not based on age or background, but on the actual lived experience of the user. Finally, the visitors/residents continuum makes no value judgments on its subjects. There is no presumption that a digital native would be more effective at accomplishing a particular task than a visitor, or vice versa. It is merely a way of characterizing users' online engagement.

The context distinction of this typology is particularly valuable with respect to academic reading. Readers' engagement with technology and the Web in their personal lives may be vastly different from that in their professional or academic lives. Studies that explore students' current practices show that digital devices are being integrated into reading workflows to great success at the individual level, but the process of integration is unknown, and not all devices a student has access to are being used to their full potential. For instance, Kennedy et al. described a general survey of learning technology owned and used by Australian first year undergraduates in which they summarize: "the transfer from a social or entertainment technology (a living technology) to a learning technology is neither automatic nor guaranteed." [32]

4.2 Technology Imposition and Emergence

In many fields, the introduction of a superior technology signals the imminent demise of the preceding standard. Notable examples such as the punch card, the fax machine, and the public telephone booth, have been, through no fault of their own, replaced by a superior version or simply made obsolete by shifting behaviour patterns. Other fields and technologies, such as the radio, evolve and adapt to change and become a necessary stepping stone towards the creatinon of their successors.

When it comes to reading, we have observed close to two decades of predicting the death or obsolescence of the book. Since the appearance of digital technology, theorists have been eagerly anticipating or cowering before the eventual extinction of the printed document. The preconditions are there: technology capable of instantly finding, downloading and storing hundreds of printed volumes, synchronize them across devices and the Web, and digital devices becoming more ubiquitous, capable and affordable than ever before. Paper is expected to play a much smaller role in offices, in casual and leisure reading, and in the academic world.

However, despite over a decade of research and technological advancement that has seen reading devices become lighter, more affordable and more diverse in form and function, the dream of successful institution-wide deployments of academic reading devices has not been realized. Recent pilot studies of such deployments [11] [76] report similar patterns:

- e-readers offer poor support for necessary companion activities such as highlighting, annotation and multi-document navigation
- use overwhelmingly declines over the course of the deployment and once the novelty effect wears off, students fail to integrate these devices into their natural reading workflows
- reading from screens is not as easy or intuitive as reading from paper
- e-reader portability is negated by the fact that they are often used alongside other, less portable technologies.

Much of this is due to the poor support in most mainstream digital technology for the companion activities to responsive reading we discussed in Chapter 3. Some of these results may also be explained by the phenomenon of technology imposition: the notion that there is a single device or technology that can be deployed en masse to supersede not only paper, but also any other devices in the academic world. The notion is certainly appealing and welcome to large educational institutions that would like to facilitate study material distribution to their cohorts by recommending a single solution, and to device manufacturers who would love nothing more than to position their device as *the* alternative to paper chosen by the institutions in question. The reality is that every single widespread deployment or pilot project of a single digital reading device has failed to gain traction among students. Young [93] described the results of attempting a digital textbook lending pilot project with a basic e-reader. Students overwhelmingly rejected the devices, citing battery life and navigation issues, and reverted to printed textbooks. Behler [11] deployed e-readers to over two hundred university-level readers in first year seminar classrooms and through a library lending program, and received critical feedback on the support for companion activities such as annotation. Thayer et al. [76] conducted a long-term deployment that was provided an 11 inch e-reader as an alternative device to students' usual technologies. Although initially enthusiastic, participants in the study quickly phased the device out of their reading routines.

The studies mentioned above all cite a combination of companion activity support issues and individual preferences as the reasons for the lack of adoption. This is strong evidence against attempting what Thayer et al. [76] called the imposition of digital reading technology.

Reading, like learning, is a highly individualized activity and the creation and deployment of a single solution to "solve" digital reading is highly unlikely, for two crucial reasons. The first is that academic readers' needs are highly variable, not only by field and level of study but also by individual learning and reading styles, location and even time of day. The second reason is that the current patterns of technology acquisition and use hint at a shift away from dedicated specialized devices and to an ecosystem of connected, somewhat interchangeable devices with the capability to participate in multi-device shared workflows for individual users. This trend makes it highly unlikely that university students who already carry two or three general purpose digital devices with them (e.g., netbook, smartphone, tablet), would ever adopt and consistently use a single academic reading device, regardless of how well suited it was to the task.

Given the highly individual nature of reading, it is to be expected that every academic reader will have a different set of strategies and technologies that he or she utilizes to successfully complete reading tasks for school. Rose [64] found that students make a conscious and conscientious effort to integrate digital devices into their study practices. In the absence of a widespread standard for such integration, students are using different technologies (including paper) according to each one's particular strengths and the availability of the technology to that particular student. The result is an emergent successful integration of heterogeneous devices into coherent workflows that are catered both to the individuals' needs and to the devices they use every day.

We believe that rather than conducting deployment studies with insular technology, research efforts should focus on harnessing the creative power of the millions of successful academic readers who have come up with unique, customized reading workflows. By researching how academic readers use the technologies available to them, and where they are hindered by technological limitations, it is possible to facilitate the integration and coexistence of the technologies already available to students into one coherent ecosystem of devices.

4.3 Integration and Multi-device Workflows

With the advancement of multiple connected devices, their interaction and coexistence have become the subject of much research, with most of it having been conducted on early adopters as the only viable population with access to such a range of devices. As more general populations such as professionals and students adopt multi-device workflows, research can focus on, and generalize from, their experience.

In the context of navigation, we posited that multi-document use cases should be regarded as the default for responsive reading by knowledge workers and academics. Similarly, we believe that responsive reading workflows should also be assumed to involve not only multiple documents, but also multiple devices. This is due not only to multiple device ownership among knowledge workers, but also due to the overwhelming evidence that there is no single device that can reliably and easily support all of the companion activities detailed in Chapter 3. Given that different devices are best suited to different kinds of content consumption and creation, a far easier and likelier workflow involves the emergent use of each device for its strengths. If that is the case, then the ways in which devices *interact* with each other becomes at least as important as the properties of each individual device.

4.3.1 Knowledge Workflows

Santosa and Wigdor [67] conducted a study on a sample of twenty-two knowledge workers across disciplines. They were found to manage an average of about ten devices each, including around three PC's and three portable devices (e.g., tablet, smartphone, e-reader). It is important to note that these devices are very rarely homogeneous, both in form factor and in platform or manufacturer. A worker may own a desktop PC for computing power and a laptop for work on the road, and perhaps a netbook for entertainment, but she is unlikely to have five identical laptops. These devices are all different, some are highly specialized, and they could decidedly benefit from synergistic use.

This study investigated workflows of content consumption and creation in a professional capacity. It detailed four patterns of use, which included serial and simultaneous (parallel) use of devices (digital or non-digital).

- **Producer/Consumer** (serial): Users found or created a document on one device, then transferred it to different device for reading or review. For students, this would involve downloading an article or lecture using one's PC, then transferring it to a tablet to read.

- **Performer/Informer** (parallel): One device contained the main content creation task while another was used as an information access portal. Students would use one device to write a paper and one to consult references.

- **Performer/Informer/Helper** (parallel): A third small device was used to perform calculations, cross off tasks, etc. alongside information access portal and main device. Typically one device is involved in content creation, one in research and content consumption and one is a secondary supporting device.

- **Controller-Viewer/Analyzer** (parallel): Combining multiple devices in the service of one task. Examples include a remote and a display, or a soft keyboard and a typing surface, on different devices.

This patterns observed in this study are highly applicable to academics and to students who are potentially using multiple devices in their studying workflows. Santosa and Wigdor [67] also noted the role of non-digital artifacts such as scrap paper for notetaking (observed in 18 of 22 participants), printed documents consulted as part of a Performer/Informer pattern and whiteboards for collaborative work and task lists. These findings are important to understanding device integration and should be applied more specifically to the domain of academic reading. In the next chapter, we describe a study we are conducting that focuses on the task of responsive reading in the context of academics and knowledge workers in an educational setting.

Chapter 5

Future Work and Research Directions

After an overview of the nature and cognition of reading, the types of reading technologies available, the requirements and companion activities of responsive reading and the approaches to studying theses activities, we have determined that the understanding of heterogeneous multi-device and multi-technology workflows in academics will be necessary to the creation and testing of any academic reading solutions that embrace and account for the emergent nature of reading habits and readers' existing competencies.

5.1 Exploratory Study of Academic Reading

We have devised an exploratory study on the multi-technology workflows of students and postgradutes reading documents for academic purposes. Studies that explore students' current practices show that digital devices are being used in responsive reading to great success at the individual level [32], but the process of heterogeneous device integration is unknown, and not all devices a student has access to are being used to their full potential.

The primary research question motivating this study is: "How are academic readers (students and postgraduates) managing required reading tasks in the context of heterogeneous ecosystems of reading technologies?". Specific subquestions concern what devices are available to users but are not being used for reading tasks, what users perceive are bottlenecks, gaps or failures in functionality, the capabilities of the ecosystem or their own satisfaction. One of the important results we hope to obtain is some set of common patterns or recommendations that can guide future work to mitigate or address these moments of failure.

We will investigate how students are using multiple technologies, including paper, to best support their academic reading. We will also probe for issues of data fragmentation: how students manage and convert digital and non-digital material and annotations, open and proprietary file formats and whether they find, as noted by Santosa and Wigdor [67] for knowledge workers, that their data is sometimes "trapped" on specialized devices.

The study will consist of three stages.

- **Stage I: Questionnaire**
  Participants will complete a basic demographics questionnaire, a short reading comprehension task, an unaltered technology self-efficacy scale [55] supplemented with online reading-specific questions which will be analyzed separately, and a survey of technologies participants own and their perceived frequency of use in academic reading contexts.

- **Stage II: Data Collection**
  Participants will be asked to report on their academic reading activities over the course of a month. Each entry is meant to encompass one reading session, i.e. a period of 10 or more minutes reading in the same location. It may involve multiple documents and multiple devices. Participants will be asked to report on the number and type of documents read, content created, and any issues they had with the workflow. They will also be instructed to take pictures of their reading layouts, which will be submitted alongside their diary entries. Twice-weekly e-mail reminders will be sent to participants to encourage them to report eligible sessions.

- **Stage III: Interviews**
  Data from the diary study will be used to prompt discussion of reading practices, strategies and shortcomings in one-hour semi-structured individual interviews. These interviews will serve to obtain nuanced and detailed data about participants' reading experiences over the course of the month and allow them to voice any systemic observations or concerns.

This study has received ERB approval and will be completed in April 2014. The complete ethics proposal is available as Appendix A at the end of this document.

5.2 Research Directions and Future Work

We hypothesize that the results of the exploratory reading study will indicate a need to facilitate data exchange and simultaneous management of heterogeneous devices to support academic reading. Based on these results, a set of design considerations and requirements will be compiled, and a research direction will be decided upon, ideally by May 2014.

A significant subset of these considerations will be implemented into a prototype integration platform, which would likely be geared specifically towards academic reading. The main goals of the platform would be cross-platform integration and non-prescriptive workflows. It will provide users with additional tools to enable multiple device interaction and to support the full range of companion activities to responsive reading, but it will be first and foremost adaptable to readers' existing habits and competencies and highly customizable. The platform can be completed by the end of 2014 and may build upon Wigdor's Symphony of Devices concept and screen management framework.

An evaluation of the platform with students would follow, emphasizing the principal areas in which the intervention is designed to have an impact. Possible measures include, but not limited to, reader satisfaction, reading comprehension, ease and consistency of device incorporation, and, time permitting, medium- to long-term end user adoption.

Bibliography

[1] Adler, A., Gujar, A., Harrison, B. L., O'hara, K., and Sellen, A. A diary study of work-related reading: design implications for digital reading devices. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (1998), ACM Press/Addison-Wesley Publishing Co., pp. 241–248.

[2] Adler, M. J., and Van Doren, C. L. How to read a book, 1972.

[3] Agarawala, A., and Balakrishnan, R. Keepin'it real: pushing the desktop metaphor with physics, piles and the pen. In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (2006), ACM, pp. 1283–1292.

[4] Anderson, R. C., and Pearson, P. D. A schema-theoretic view of basic processes in reading comprehension. *Handbook of reading research 1* (1984), 255–291.

[5] Arthur, R., and Olsen Jr, D. R. Xice windowing toolkit: Seamless display annexation. *ACM Transactions on Computer-Human Interaction (TOCHI) 18*, 3 (2011), 14.

[6] Baddeley, A., Della Sala, S., Robbins, T., and Baddeley, A. Working memory and executive control [and discussion]. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences 351*, 1346 (1996), 1397–1404.

[7] Barrett, E. *Text, ConText, and HyperText: Writing with and for the Computer*. MIT Press, 1988.

[8] Barrouillet, P., Bernardin, S., and Camos, V. Time constraints and resource sharing in adults' working memory spans. *Journal of Experimental Psychology: General 133*, 1 (2004), 83.

[9] Baudrillard, J., and Turner, C. *Cool Memories II, 1987-1990*. Duke University Press (Durham), 1996.

[10] Beeman, W. O., Anderson, K. T., Bader, G., Larkin, J., McClard, A. P., McQuillan, P., and Shields, M. Hypertext and pluralism: from lineal to non-lineal thinking. In *Proceedings of the ACM conference on Hypertext* (1987), ACM, pp. 67–88.

[11] Behler, A. E-readers in action. *American Libraries 40*, 10 (2009), 56–59.

[12] Berners-Lee, T., and Cailliau, R. Worldwideweb: Proposal for a hypertext project. *Retrieved on February 26* (1990), 2008.

[13] Brown, A. L., and Day, J. D. Macrorules for summarizing texts: The development of expertise. *Journal of verbal learning and verbal behavior 22*, 1 (1983), 1–14.

[14] Bush, V. As we may think. *Atlantic Monthly* (1945).

[15] Chen, N., Guimbretiere, F., Dixon, M., Lewis, C., and Agrawala, M. Navigation techniques for dual-display e-book readers. In *Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems* (2008), ACM, pp. 1779–1788.

[16] Chen, N., Guimbretiere, F., and Sellen, A. Designing a multi-slate reading environment to support active reading activities. *ACM Transactions on Computer-Human Interaction (TOCHI) 19*, 3 (2012), 18.

[17] Chen, N., Guimbretière, F., and Sellen, A. Graduate student use of a multi-slate reading system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2013), ACM, pp. 1799–1808.

[18] Cicerón, M. T. *Cicero: On the Ideal Orator (De Oratore).* Oxford University Press, 2001.

[19] Cowan, N. *Attention and memory.* Oxford University Press, 1995.

[20] Crane, G., and Mylonas, E. Ancient materials, modern media: Shaping the study of classics with hypertext. In *Hypermedia and literary studies* (1991), MIT Press, pp. 205–220.

[21] Duke, N. K., and Pearson, P. D. Effective practices for developing reading comprehension. *What research has to say about reading instruction 3* (2002), 205–242.

[22] Elfenbein, A. Cognitive science and the history of reading. *PMLA* (2006), 484–502.

[23] Ericsson, K. A., and Kintsch, W. Long-term working memory. *Psychological review 102*, 2 (1995), 211.

[24] Fitzgibbons, M. Implications of hypertext theory for the reading, organization, and retrieval of information. *Library Philosophy and Practice (e-journal)* (2008), 170.

[25] Flavell, J. H. Cognitive monitoring. *Children's oral communication skills* (1981), 35–60.

[26] Foucault, M. The archaeology of knowledge and the discourse on language, translated by a. m. sheridan smith. *New York: Pantheon* (1972).

[27] Garner, R. *Metacognition and reading comprehension.* Ablex Publishing, 1987.

[28] Hamilton, P., and Wigdor, D. Conductor: Enabling and understanding cross-device interaction. In *CHI 2014 (Submitted for review)* (2014), ACM.

[29] Harris, J. Text annotation and underlining as metacognitive strategies to improve comprehension and retention of expository text.

[30] Hulme, C., Roodenrys, S., Brown, G., and Mercer, R. The role of long-term memory mechanisms in memory span. *British Journal of Psychology 86*, 4 (1995), 527–536.

[31] Just, M. A., and Carpenter, P. A. A theory of reading: From eye fixations to comprehension. *Psychological review 87* (1980), 329–354.

[32] Kennedy, G. E., Judd, T. S., Churchward, A., Gray, K., and Krause, K. First year students' experiences with technology: Are they really digital natives. *Australasian Journal of Educational Technology 24*, 1 (2008), 108–122.

[33] Kintsch, W. *Comprehension: A paradigm for cognition.* Cambridge university press, 1998.

[34] Kitchin, R. M. Cognitive maps: What are they and why study them? *Journal of environmental psychology 14*, 1 (1994), 1–19.

[35] Knight, K., and Marcu, D. Summarization beyond sentence extraction: A probabilistic approach to sentence compression. *Artificial Intelligence 139*, 1 (2002), 91–107.

[36] Kol, S., and Schcolnik, M. Reading from screen vs. reading from paper: A pilot study. *CAELL Journal 8*, 1 (1997), 10–14.

[37] Kolers, P. A. Three stages of reading. *Basic studies on reading. New York: Basic Books* (1970), 90–118.

[38] Levin, J. R., Anglin, G. J., and Carney, R. N. On empirically validating functions of pictures in prose. *The psychology of illustration 1* (1987), 51–85.

[39] Liu, Z. Reading behavior in the digital environment: Changes in reading behavior over the past ten years. *Journal of Documentation 61*, 6 (2005), 700–712.

[40] Lorch, R. F., Lorch, E. P., and Klusewitz, M. A. College students' conditional knowledge about reading. *Journal of educational psychology 85*, 2 (1993), 239.

[41] Malone, T. W. How do people organize their desks?: Implications for the design of office information systems. *ACM Transactions on Information Systems (TOIS) 1*, 1 (1983), 99–112.

[42] Mani, I., and Maybury, M. T. *Advances in automatic text summarization.* the MIT Press, 1999.

[43] Marshall, C. C. Annotation: from paper books to the digital library. In *Proceedings of the second ACM international conference on Digital libraries* (1997), ACM, pp. 131–140.

[44] Marshall, C. C. Reading and writing the electronic book. *Synthesis Lectures on Information Concepts, Retrieval, and Services 1*, 1 (2009), 1–185.

[45] Marshall, C. C., and Bly, S. Turning the page on navigation. In *Digital Libraries, 2005. JCDL'05. Proceedings of the 5th ACM/IEEE-CS Joint Conference on* (2005), IEEE, pp. 225–234.

[46] Marshall, C. C., and Brush, A. B. Exploring the relationship between personal and public annotations. In *Digital Libraries, 2004. Proceedings of the 2004 Joint ACM/IEEE Conference on* (2004), IEEE, pp. 349–357.

[47] Marshall, C. C., Price, M. N., Golovchinsky, G., and Schilit, B. N. Introducing a digital library reading appliance into a reading group. In *Proceedings of the fourth ACM conference on Digital libraries* (1999), ACM, pp. 77–84.

[48] Marshall, C. C., Price, M. N., Golovchinsky, G., and Schilit, B. N. Designing e-books for legal research. In *Proceedings of the 1st ACM/IEEE-CS joint conference on Digital libraries* (2001), ACM, pp. 41–48.

[49] Marshall, C. C., and Ruotolo, C. Reading-in-the-small: a study of reading on small form factor devices. In *Proceedings of the 2nd ACM/IEEE-CS joint conference on Digital libraries* (2002), ACM, pp. 56–64.

[50] Marshall, C. C., and Shipman III, F. M. Spatial hypertext and the practice of information triage. In *Proceedings of the eighth ACM conference on Hypertext* (1997), ACM, pp. 124–133.

[51] McDermott, M. Knowledge workers: How can you gauge their effectiveness? *Leadership Excellence 22*, 10 (2005), 15.

[52] McKnight, C. *Hypertext in context*, vol. 6. Cambridge University Press, 1991.

[53] Mento, A. J., Martinelli, P., and Jones, R. M. Mind mapping in executive education: applications and outcomes. *Journal of Management Development 18*, 4 (1999), 390–416.

[54] Miller, G. A. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological review 63*, 2 (1956), 81.

[55] Miltiadou, M., and Yu, C. H. *Validation of the online technologies self-efficacy scale (OTSES).* ERIC Clearinghouse, 2000.

[56] Morris, M. R., Brush, A. B., and Meyers, B. R. Reading revisited: Evaluating the usability of digital display surfaces for active reading tasks. In *Horizontal Interactive Human-Computer Systems, 2007. TABLETOP'07. Second Annual IEEE International Workshop on* (2007), IEEE, pp. 79–86.

[57] Nelson, T. H. The hypertext. *proceedings of the World Documentation Federation* (1965), 84–100.

[58] O'Hara, K. Towards a typology of reading goals. Tech. rep., 1996.

[59] O'Hara, K., and Sellen, A. A Comparison of Reading Paper and On-Line Documents. In *the SIGCHI conference* (New York, New York, USA, 1997), ACM Press, pp. 335–342.

[60] O'HARA, K. P., Taylor, A., Newman, W., and Sellen, A. J. Understanding the materiality of writing from multiple sources. *International journal of human-computer studies 56*, 3 (2002), 269–305.

[61] Prensky, M. Digital natives, digital immigrants part 1. *On the horizon 9*, 5 (2001), 1–6.

[62] Protopsaltis, A., and Bouki, V. Towards a hypertext reading/comprehension model. In *Proceedings of the 23rd annual international conference on Design of communication: documenting & designing for pervasive information* (2005), ACM, pp. 159–166.

[63] Pugh, A. K. Silent reading. *An introduction to its study and teaching. London: Heinemann* (1978).

[64] Rose, E. The phenomenology of on-screen reading: University students' lived experience of digitised text. *British Journal of Educational Technology 42*, 3 (2011), 515–526.

[65] Rothkopf, E. Z. Incidental memory for location of information in text. *Journal of Verbal Learning and Verbal Behavior 10*, 6 (1971), 608–613.

[66] RUMELHART, D. E. *Schemata: The building blocks of cognition.* Center for Human Information Processing, University of California, San Diego, 1978.

[67] SANTOSA, S., AND WIGDOR, D. A field study of multi-device workflows in distributed workspaces. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing* (2013), ACM, pp. 63–72.

[68] SCHILIT, B. N., GOLOVCHINSKY, G., AND PRICE, M. N. Beyond paper: supporting active reading with free form digital ink annotations. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (1998), ACM Press/Addison-Wesley Publishing Co., pp. 249–256.

[69] SELLEN, A. J., AND HARPER, R. H. *The myth of the paperless office.* The MIT Press, 2003.

[70] STANOVICH, K. E. Toward an interactive-compensatory model of individual differences in the development of reading fluency. *Reading research quarterly* (1980), 32–71.

[71] STEIN, R. Page versus pixel. *Feed Magazine* (1995).

[72] TAJIKA, T., YONEZAWA, T., AND MITSUNAGA, N. Intuitive page-turning interface of e-books on flexible e-paper based on user studies. In *Proceedings of the 16th ACM international conference on Multimedia* (2008), ACM, pp. 793–796.

[73] TAN, D. S., PAUSCH, R., STEFANUCCI, J. K., AND PROFFITT, D. R. Kinesthetic cues aid spatial memory. In *CHI'02 extended abstracts on Human factors in computing systems* (2002), ACM, pp. 806–807.

[74] TASHMAN, C. S., AND EDWARDS, W. K. Active reading and its discontents: the situations, problems and ideas of readers. In *Proceedings of the 2011 annual conference on Human factors in computing systems* (2011), ACM, pp. 2927–2936.

[75] TASHMAN, C. S., AND EDWARDS, W. K. Liquidtext: a flexible, multitouch environment to support active reading. In *Proceedings of the 2011 annual conference on Human factors in computing systems* (2011), ACM, pp. 3285–3294.

[76] THAYER, A., LEE, C. P., HWANG, L. H., SALES, H., SEN, P., AND DALAL, N. The imposition and superimposition of digital reading technology: the academic potential of e-readers. In *CHI '11: Proceedings of the 2011 annual conference on Human factors in computing systems* (May 2011), ACM.

[77] THINYANE, H., AND THORNE, J.-P. An empirical investigation into the use of digital photo frames as low cost e-book readers. In *Proceedings of the South African Institute for Computer Scientists and Information Technologists Conference* (2012), ACM, pp. 270–276.

[78] THORNGATE, W. On paying attention. *Recent trends in theoretical psychology* (1988), 247–264.

[79] TOLMAN, E. C. Cognitive maps in rats and men. *Psychological review 55*, 4 (1948), 189.

[80] TULVING, E., AND GOLD, C. Stimulus information and contextual information as determinants of tachistoscopic recognition of words. *Journal of Experimental Psychology 66*, 4 (1963), 319.

[81] van den Broek, P., Gustafson, M., Goldman, S., Graesser, A., and van den Broek, P. Comprehension and memory for texts: Three generations of reading research. *Narrative comprehension, causality, and coherence: Essays in honor of Tom Trabasso* (1999), 15–34.

[82] van den Broek, P., Kremer, K. E., et al. The mind in action: What it means to comprehend during reading. *Reading for meaning: Fostering comprehension in the middle grades* (2000), 1–31.

[83] Van den Broek, P., Lorch, R. F., Linderholm, T., and Gustafson, M. The effects of readers' goals on inference generation and memory for texts. *Memory & Cognition 29*, 8 (2001), 1081–1087.

[84] Van den Broek, P., Young, M., Tzeng, Y., Linderholm, T., van Oostendorp, H., and Goldman, S. The landscape model of reading: Inferences and the online construction of a memory representation. *The construction of mental representations during reading* (1999), 71–98.

[85] Wallen, E., Plass, J. L., and Brünken, R. The function of annotations in the comprehension of scientific texts: Cognitive load effects and the impact of verbal ability. *Educational Technology Research and Development 53*, 3 (2005), 59–71.

[86] Waller, R. What electronic books will have to be better than. *Information Design Journal 5*, 1 (1986), 72–75.

[87] White, D. S., and Le Cornu, A. Visitors and residents: A new typology for online engagement. *First Monday 16*, 9 (2011).

[88] Wilson, R., and Landoni, M. Evaluating the usability of portable electronic books. In *Proceedings of the 2003 ACM symposium on Applied computing* (2003), ACM, pp. 564–568.

[89] Wolfe, J. L. Effects of annotations on student readers and writers. In *Proceedings of the fifth ACM conference on Digital libraries* (2000), ACM, pp. 19–26.

[90] Wood, R., Baxter, P., and Belpaeme, T. A review of long-term memory in natural and synthetic systems. *Adaptive Behavior 20*, 2 (2012), 81–103.

[91] Yang, J., and Wigdor, D. Panelrama: Enabling easy specification of cross-device web applications. In *CHI 2014 (Submitted for review)* (2014), ACM.

[92] Yi, W., Park, E., and Cho, K. E-book readability, comprehensibility and satisfaction. In *Proceedings of the 5th International Conference on Ubiquitous Information Management and Communication* (2011), ACM, p. 38.

[93] Young, J. R. 6 lessons one campus learned about e-textbooks. *Chronicle of Higher Education 55*, 39 (2009).