



Passages: Interacting with Text Across Documents

Han Han, Junhang Yu, Raphael Bournet, Alexandre Ciorascu, Wendy E. Mackay, Michel Beaudouin-Lafon

► To cite this version:

Han Han, Junhang Yu, Raphael Bournet, Alexandre Ciorascu, Wendy E. Mackay, et al.. Passages: Interacting with Text Across Documents. CHI 2022 - ACM Conference on Human Factors in Computing Systems, Apr 2022, New Orleans, LA, United States. pp.1-17, 10.1145/3491102.3502052 . hal-03664173

HAL Id: hal-03664173

<https://hal.science/hal-03664173>

Submitted on 10 May 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Passages: Interacting with Text Across Documents

Han L. Han

han.han@lisn.fr

Université Paris-Saclay, CNRS, Inria
Laboratoire Interdisciplinaire en
Sciences du Numérique (LISN)
Orsay, France

Junhang Yu

junhang.yu@upsaclay.fr

Université Paris-Saclay, CNRS, Inria
Laboratoire Interdisciplinaire en
Sciences du Numérique (LISN)
Orsay, France

Raphael Bournet

raphael.bournet@upsaclay.fr

Université Paris-Saclay, CNRS, Inria
Laboratoire Interdisciplinaire en
Sciences du Numérique (LISN)
Orsay, France

Alexandre Ciorascu

alexandre.ciorascu@upsaclay.fr

Université Paris-Saclay, CNRS, Inria
Laboratoire Interdisciplinaire en
Sciences du Numérique (LISN)
Orsay, France

Wendy E. Mackay

mackay@lisn.fr

Université Paris-Saclay, CNRS, Inria
Laboratoire Interdisciplinaire en
Sciences du Numérique (LISN)
Orsay, France

Michel Beaudouin-Lafon

mbl@lisn.fr

Université Paris-Saclay, CNRS, Inria
Laboratoire Interdisciplinaire en
Sciences du Numérique (LISN)
Orsay, France

ABSTRACT

A key aspect of knowledge work is the analysis and manipulation of sets of related documents. We conducted interviews with 12 patent examiners and 12 scientists and found that all face difficulties using specialized tools for managing text from multiple documents across interconnected activities, including searching, collecting, annotating, organizing, writing and reviewing, while manually tracking their provenance. We introduce *Passages*, interactive objects that reify text selections and can then be manipulated, reused, and shared across multiple tools. *Passages* directly supports the above-listed activities as well as fluid transitions among them. Two user studies show that participants found *Passages* both elegant and powerful, facilitating their work practices and enabling greater reuse and novel strategies for analyzing and composing documents. We argue that *Passages* offers a general approach applicable to a wide variety of text-based interactions.

CCS CONCEPTS

• **Human-centered computing** → HCI theory, concepts and models.

KEYWORDS

Knowledge Work, Document Management, Provenance, Sensemaking, Active Reading, Reification

ACM Reference Format:

Han L. Han, Junhang Yu, Raphael Bournet, Alexandre Ciorascu, Wendy E. Mackay, and Michel Beaudouin-Lafon. 2022. *Passages: Interacting with Text Across Documents*. In *CHI 2022: Human Factors in Computing Systems May 2022, New Orleans, LA, USA*. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3491102.3502052>

1 INTRODUCTION

We are interested in developing tools for supporting knowledge work in an age of information overload [57, 58, 92]. Studies of knowledge work [3, 63, 64] highlight the complexity of this process, which involves active reading, search, retrieval, annotation and writing activities. Our goal is to gain a better understanding not only of how knowledge workers find and make sense of documents, but also how they keep track of these documents across applications, with the goal of improving tools to support their work.

Many professions have developed specialized productivity software [65, 98] to support this process. For example, legal professionals take advantage of specialized online services to find and cite relevant law books and articles that serve as precedents for their cases [63]. Scientists follow a similar process when searching for related articles that support their arguments, using the research literature to “describe, summarize, evaluate, clarify and/or integrate the content of primary reports” [19].

Unfortunately, current software typically traps information into *information silos* [46, 71], making it difficult to reuse content or obtain a unified view of the material [65]. This poses a serious design challenge, since adding yet another tool risks complicating, rather than supporting, the complex, multi-faceted nature of document management. The first research question (RQ1) thus focuses on better understanding how today’s knowledge workers currently perform active reading in their document management process, with particular emphasis on their use of existing software tools.

A second key problem knowledge workers face is capturing and maintaining *provenance* [18]—tracking the source of a document and returning to it—despite varied formats and diverse, non-interconnected sources. Evans et al. [22] argue that tracking provenance is a key factor in the media industry’s current “crisis in journalism” [82], and Flinham et al. [25] show how identifying an article’s source is essential for evaluating its trustworthiness. Similarly, scientists rely on provenance to ensure reproducibility and uncover potential plagiarism [18], and Jensen et al.’s [44] longitudinal study of knowledge workers finds that provenance cues significantly aid recall. The second research question (RQ2) thus asks how today’s knowledge workers manage provenance.

This paper first reviews related research and then describes two interview studies, which helped us gain a detailed understanding of current document management practices. We chose two different types of “lead users” [88]: the highly codified profession of patent examiners, and the more open-ended profession of research scientists. We then introduce the concept of *Passages*, illustrate it with a scenario, and describe a proof-of-concept implementation. Next, we describe the results of a structured observation with 12 scientists and a generative walkthrough with six participants from a patent organization. We conclude with directions for future research.

2 RELATED WORK

Knowledge workers engage in *active reading* of documents, both to make sense of them and to evaluate each document’s relevance to their needs. When they find useful documents, they must also keep track of their sources, or maintain the provenance of each document, in order to cite them accurately and return to them if needed. We thus review related work with respect to three key areas within knowledge work: active reading of documents, sensemaking, and information provenance.

2.1 Active Reading of Documents

Practices. Knowledge workers engage in “active reading” of documents, which involves interweaving reading with a variety of associated activities [3, 29, 30]. For example, O’Hara et al. [64] interviewed researchers about their library use and characterize scholarly research as “a complex process of searching, information retrieval, reading, information extraction and recording by annotation and note-taking, information review, and writing new compositions”. In a diary study of knowledge workers’ reading practices, Adler et al. [3] found that reading goes hand-in-hand with writing, and that knowledge workers often read multiple documents in parallel, a finding echoed by Tashman and Edwards’s study of active reading [85]. Marshall et al. [63] found that law students switch frequently between annotating, organizing and writing. Although this research shows that active reading includes a diverse set of activities, we do not fully understand how current software tools support transitions among them.

Systems. Multiple systems support document reading, including active reading [40, 63, 75, 86], active diagramming [84] and active note-taking [41], as well as more specific tasks such as annotation [73, 95, 96] and navigation [4, 93]. Several systems explicitly support document-reading practices, such as *InkSeine* [41], which accommodates searching during active note taking. By transforming search queries into first class objects, users can quickly capture and save search results into their “ink” notebooks. *GatherReader* [40] builds on the finding that reading co-occurs with writing and cross-referencing documents [3], as well as gathering pieces of information [61]. It supports collecting multiple objects via a temporal visual clipboard, using pen and touch interaction. *LiquidText* [86] offers readers highly flexible and malleable documents, with fluid representations and a multi-touch gesture-based interface. These systems focus on reading-related activities, particularly note-taking [40, 86] and searching [30, 41]. However, we are also interested in helping readers manage active reading while dealing with multiple document-based applications.

2.2 Sensemaking

Practices. Sensemaking is defined as creating a mental representation of an information space to support the user’s goals [74]. Sensemaking research explores relationships across different knowledge work activities, especially searching, capturing and organizing information. Kittur et al. [49] differentiate between two main phases—first information seeking then sensemaking—and highlight the cost of creating structure too early in the foraging process.

Other researchers focus on distributed [24] and collaborative [66] sensemaking, or the role of resources [78]. Sellen et al. [77] identify six categories of knowledge workers’ use of the web, and found that most searches lead to further activities, especially referring back to a result and incorporating it into a document. They argue that users need more flexible ways of saving text and search results.

Other studies of web-based information seeking highlight users’ revisitation behavior [2, 55] and their use of multiple windows [21, 89, 90], as well as activities beyond searching [14, 77]. A key sensemaking activity involves creating external representations, such as tables [54], hierarchies [1], diagrams [84], networks [35], canvases [51], and more [79]. We are interested in the knowledge worker’s complete workflow, encompassing the interconnection of active reading and sensemaking activities.

Systems. A number of systems support sensemaking, including searching [38, 60, 70, 91], collecting [42, 49] and organizing [15, 35, 62, 72] information, as well as systems for creating tables [68], graphs and hypermedia structures [74]. *Hunter Gatherer* [76], *Unakite* [54], *ScratchPad* [32], and Dontcheva et al.’s web summarization system [20] support searching, collecting and organizing information, but do not extend to writing or reviewing final documents, e.g. summaries and analysis reports [85], which are also a part of knowledge work [64]. Other systems such as *CiteSense* [99] and *Entity Workspace* [10] provide an integrated environment with a multi-panel interface, but do not support the flexibility needed for the diverse and changing nature of knowledge work [48, 85]. More generally, such systems are not designed to support the interconnected activities required for reading and sensemaking.

2.3 Provenance

Pérez et al. [69] define *provenance* as: “the entire amount of information, comprising all the elements and their relationships, that contribute to the existence of a piece of data.” Jensen et al.’s [44] investigation of the provenance of files on knowledge workers’ desktops shows how it reveals their work patterns. Several systems accommodate file provenance, e.g. for version management [47, 53] and file retrieval [28, 81, 83]. *TaskTrail* [83] tracks file provenance from copy-paste and save-as commands to help users re-find documents. *FileWeaver* [31] tracks file dependencies and can recreate dependent files when a source is modified. Karlson et al.’s [47] system tracks copy relationships among files to help users manage versions. Given the highly networked character of today’s information systems, Lindley et al. [53] challenge the original file metaphor [80]. They introduce “file biographies” to capture the provenance of each file and let users track how it propagates. Despite research showing how knowledge workers often focus on short *snippets* of information [54, 76, 84, 85], none of these systems directly support the capture and maintenance of *snippet* provenance.

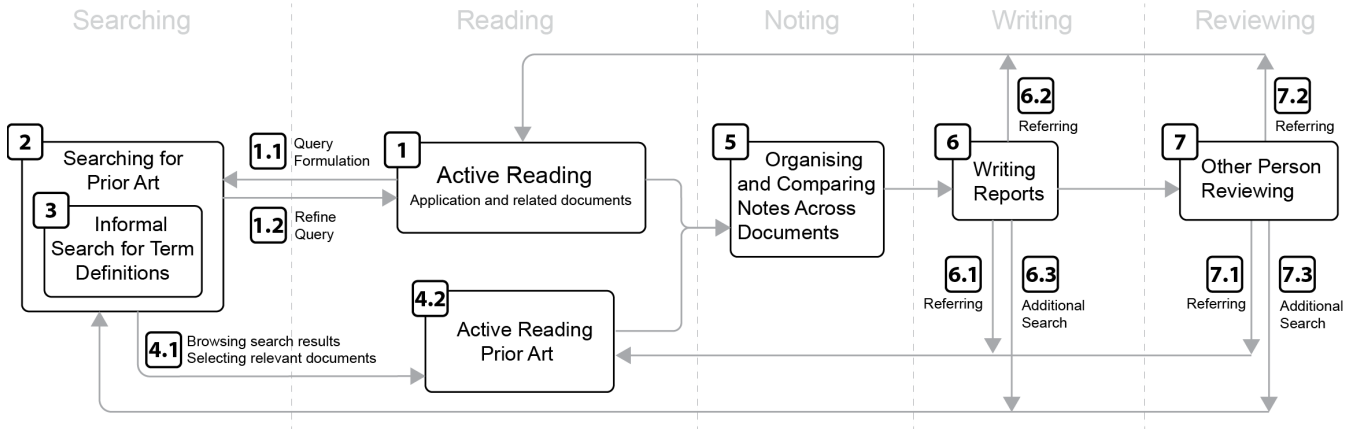


Figure 1: Patent examiners engage in a series of document-related activities when searching for prior art, including formulating search terms; reading and annotating documents; and writing and reviewing their own and other examiners’ reports.

Although these systems each offer useful support for different aspects of knowledge work, including active reading, sensemaking and determining provenance, we are interested in designing software tools that work *across* these activities. We decided to examine two groups of knowledge workers, patent examiners and research scientists, who read, interpret, organize and share documents across multiple applications over long periods of time, to contribute insights to the design of such software tools.

3 STUDY 1: INTERVIEWS WITH PATENT EXAMINERS

The legal profession offers an extreme example of document-intensive knowledge work, with numerous design challenges [3, 63]. Our earlier work focused on lawyers and patent writers [36]. Here, we explore the issues faced by a related group—*patent examiners*—who must search for, analyze and communicate effectively about complex legal documents in order to make critical decisions about intellectual property. To address RQ1, we decided to interview professional patent examiners to better understand the document-related aspects of the patent examination process, as well as their current use of software tools. We also asked how they capture and maintain the provenance of their documents (RQ2).

3.1 Method¹

Participants. We recruited 12 participants (11 men, 1 woman; mean=13 years of experience), from a major patent organization. Nine are full-time patent examiners, and three are part-time patent examiners—two software engineers and a product owner from the software development team. All use the organization’s internal software applications to search and examine patent applications, and each specializes in a particular technical field, including antennas, biotechnology, CPUs, displays, medicine, optics, and polymers.

Setup. All interviews were conducted in English via video calls.

Procedure. Each interview lasted two-three hours. In addition to collecting basic information about each patent examiner’s background, we conducted story interviews [56] where they were asked to describe a recent, memorable event related to searching and examining a specific patent application. Participants were encouraged to show the specific documents and software involved, and to re-enact each search process, step by step. We also interviewed members of the development team for a different perspective on the challenges faced by patent examiners as well as their goals for future system development.

Data Collection. All interviews were screen recorded and transcribed. We also took hand-written notes and collected examples of documents and software screens.

Analysis Method. We analyzed the interviews using reflexive thematic analysis [12]. After interviewing six participants, one researcher generated themes and codes using both a deductive (top-down) approach and an inductive (bottom-up) approach to identify new themes. Deductive themes are: process breakdowns, user workarounds and user innovations with respect to participants’ use of software tools and management of provenance. Next, two researchers worked together to reassess the themes and codes, and group them into larger categories. They discussed any disagreements and rechecked the interview transcripts to reach a shared understanding for the entire set of interviews, arriving at the final categorization after two iterations.

3.2 Results

We collected 13 unique patent examples from the 12 patent examiners. For confidentiality reasons, participants sometimes needed to hide information and would jump to another, similar example to illustrate the relevant aspect of the process, and would then return to the main example. The next sections describe the patent examiners’ current document management process, with particular emphasis on their use of software tools (RQ1); as well as their techniques for capturing and maintaining provenance (RQ2).

¹All procedures for the four studies in this paper were approved by COERLE, our organization’s Institutional Review Board (IRB).

3.2.1 Patent examination process. To award a new patent, the patent examiner engages in a complex, highly document-centric process, with six basic steps: filing, search, examination, grant award, opposition and appeal². Here, we focus on the search step, illustrated in Fig. 1, which judges the patentability, or novelty, of an application relative to all prior art. We chose this step because it involves a variety of interrelated document-based activities, and is required for all patent examiners.

Patent examiners first read the patent application to understand its claimed invention (1). If the application is re-submitted, they must also read other related documents, such as previous search reports and letters from the attorneys, to understand the full context (1). Examiners actively annotate the application and issue queries (1.1) that search for prior art (2). Because of the highly technical nature of patents, examiners must often search for new terminology and term definitions using specialized web services (3). They next browse the search results to identify (4.1) and read relevant documents (4.2).

Study participants reported that they read from 100 to 300 documents for each patent search. Examiners next select a subset of selected documents and compare them to the patent application based on specific criteria such as novelty, inventive steps and clarity (5). This process of reading and searching is highly iterative, since examiners often learn a new term from one prior art document and use it to refine subsequent queries (1.2) to search for (1.1) and read additional documents (4.2). Most participants (10/12) mentioned that they take notes to keep track of their work (5).

At the end of this process, examiners identify three to five highly relevant prior art documents. They then write a search report, drawing from their personal notes (6). The patent institute requires them to cite specific, sentence-level evidence in their report, so they often refer back to the patent application (6.2) and prior art documents (6.1). They may also search for additional documents (6.3) if they realize that something is missing. The search report is then reviewed by the chairman of the patent division (7). This review also involves frequent references to the application (7.1) and prior art documents (7.2), resulting in additional searches (7.3).

The search phase usually leads to the examination phase, at which point the patent office decides whether or not to grant the patent. Three patent examiners examine the application, one of whom maintains contact with the patent attorney. The dialogue between the examiner and the patent attorney sometimes results in modifications to the original patent application.

3.2.2 Software tool use. We identified two key challenges with respect to understanding their current use of software tools to manage the patent search process (RQ1): *Transferring information across applications is cumbersome (T1)* and *Recording personal notes risks duplicating information (T2)*.

T1: Transferring information across applications is cumbersome. The patent organization has developed over 20 specialized applications for patent examiners. Most are designed for specific activities such as task management, report creation, document viewing, and searching for prior art. Others are designed for highly specific tasks.

For example, a chemistry examiner (P7) showed us a specialized application for finding the CAS (Chemical Abstract Service) registry number when information found on a website or database entry is messy or incomplete.

Despite the availability of these applications, examiners must spend a great deal of time choosing among them and setting them up before they can perform any real work. P2 said: “*One of the problems that examiners have is that there are many different tools and there are many different times that you have to go here and there. So it is very cumbersome.*” They must also follow a preset process when using these systems. For example, P4 must open three different applications just to view a document “*so that it has the right dossier number and then it is synchronized.*” This produces many unnecessary windows that are hard to manage and clutter the examiner’s screen. Two examiners (P3, P4) developed workarounds to avoid the rigidity of this procedure-based interaction. For example, the classification tool only sends codes to three tools, none of which is the one he wants. So P3 copy-pasted the tool’s entire classification code into a Microsoft Word document, which lets him easily search for and copy the codes he needs.

T2: Recording personal notes risks duplicating information. Examiners collect information snippets and add them to their personal notes. This not only supports the current task, but also helps them in the future. P6 called these notes a “*letter to myself*”—when the application comes back for re-examination in a year, she will not have to “*spend the same amount of time again to re-familiarize [herself] with the file.*”

Some examiners structure their notes as a comparison table (P1, P6, P9), spatial canvas (P7) or as linear text (P6, P10, P11). For example, P6 created a table in Microsoft Word to compare the application and prior art documents. She considers the table to be “*a mental help to remember what the claims are about and what are the features.*” P7 pastes text snippets and screenshots into Microsoft OneNote so that he can have “*all the information at a glance*”. Other examiners avoid organizing personal notes into an intermediate structure, and instead treat the text editor where they write the final report as their note-taking space. For example, P11 keeps the claims in his head and tries “*to start the communication [with the attorney] as soon as possible*”. P6 found that managing his collected information and notes in a separate application is “*really painful...it is really nasty. It is the duplication of information...Copying is the maximum I would like to invest.*”

3.2.3 Provenance management. We also identified two additional challenges with respect to managing provenance (RQ2): *Re-finding sources is tedious (T3)* and *Search results are easily lost (T4)*.

T3: Re-finding sources is tedious. All examiners had to manually re-find text snippets from multiple sources, including prior art, application software and other related documents, in order to verify or cite them in their report. They developed diverse strategies to reduce the time involved in collecting snippets, from using the search function (P2) or manually keeping a “link” (P7). This task is even more extreme for the patent chairman, who must locate and verify many cited passages during the review process. For example, P2 must manually count line numbers: “*If I see here that this claim is supported by the description on page 5, line 15-17, I go on page 5 and have to count to the line and I find [it].*”

²Note that this practice differs between Europe (<https://www.epo.org/learning/materials/inventors-handbook/protection/patents.html>) and the United States (<https://www.uspto.gov/patents/basics/patent-process-overview>)

Examiners frequently revisit snippets of text when they edit or review their reports, but find it tedious to copy-paste them from the organization’s internal applications into their personal notes. They also have trouble maintaining links back to the original text. Although citing snippets is tedious, examiners agree that it is critical for making evidence-based decisions and essential for communicating with others. For example, P11 (a junior examiner) said: *“I have learned that, in the examination argumentation phase, it can be really helpful to cite the specific passage. If you read from the line that this functional feature is present, it does not mean that everybody else will read it. You have to explain it. If you don’t, you will get a lengthy letter back. You have to explain it anyway and this is a waste of time for both parties. This is why I am often quoting complete sentences, saying that this feature is really there, don’t come back to me and say it is not there.”*

T4: Search results are easily lost. Examiners not only search for prior art documents but also for term definitions and synonyms on the web to help them understand an unfamiliar domain. The internal search tool keeps track of the main search activity, which is the search for prior art. However, queries and results of informal searches remain ephemeral, causing repetitive re-searching for the same information. For example, after starting a search, P5 would *“go back and maybe have another look at the dictionary or the Wikipedia”* because it gives him new ideas for a search query.

3.2.4 Inductive themes. In addition to the deductive themes associated with RQ1 and RQ2, we also identified two inductive themes: *Automatic drafting ignores human expertise (T5)* and *Sharing knowledge has few rewards (T6)*.

T5: Automatic drafting ignores human expertise. The organization developed an application for automatically generating a report from a form that the examiners fill out. Although intended to improve report-writing productivity by reducing copy-pasting and re-writing, many examiners have mixed feelings about the system and are reluctant to use it. They emphasize the importance of having a human select the right text, and the need for human-to-human communication in their intellectual work. P6 said that her team dislikes it, so she dropped it: *“I did not see the benefits. I was typing as much as before and the communication is much longer because it has all these standard sentences that I don’t really need to repeat every time.”* For her, the intellectual part is about *“selecting the passages”*. P8 and P10 emphasized that the human value lies in communication and dialogue with applicants. P8 said: *“If we are moving in the direction where the communication will be automatically generated by computer, we will be sending the message that we don’t really do the intellectual work. So [the patent attorneys] will be less likely to be convinced, the same way I am less likely to be convinced.”*

T6: Sharing knowledge has few rewards. Almost half the examiners (5/12) have their own library of synonyms for their respective domains. These libraries are a form of knowledge that examiners have accumulated over many years and can be useful both to themselves and to others. Examiners share knowledge mainly through talking to people (P4, P5), exchanging Microsoft Excel or text files (P11, P12) or using an internal system (P6, P11). P6, an experienced examiner, said that if he works in an unfamiliar field, he just goes next door and asks his colleague. This is consistent with research that shows that people prefer obtaining information from other people within

an organization [39]. Two examiners (P11, P12) also mentioned exchanging their personal synonym files with colleagues.

Although the organization has an internal system for sharing synonym libraries, one participant said that examiners gradually stopped using it, since the focus on improving productivity discouraged them from taking the extra time needed to create reusable libraries. Others were willing to share their synonyms but received no benefits if they created something that exceeded their personal needs. One exception is an examiner with a huge personal synonym database—he has around 500, in contrast to his colleagues’ 20 or 30—who constantly refines and improves it, and emails it to a dozen colleagues every six months. His role is similar to the “translator” identified in Mackay’s study of exchange of customization files [57].

3.2.5 Summary. Study 1 shows the complexity of the patent examination process, and emphasizes the difficulties patent examiners encounter as they try to manage interconnected document-intensive activities through multiple specialized but separate applications (RQ1). Although each examiner establishes multiple personal strategies for organizing information, they struggle to maintain the provenance of the resulting documents (RQ2). Examiners also worry that shifting to ever-more-automated systems will reduce the human added value of their work, and they want to ensure that communication with other (human) participants remains part of the process. They feel that the current push for increased productivity without appropriate tools reduces their incentives to share knowledge, with a corresponding negative effect on the quality of their work.

4 STUDY 2: INTERVIEWS WITH SCIENTISTS

Reviewing and analyzing the scientific literature has a number of similarities to the patent review process [23]. We decided to broaden the scope of our study to include this more open-ended form of knowledge work, and thus conducted a second series of interviews with research scientists.

4.1 Method

Participants. We recruited 12 scientists (4 women, 8 men) from the following disciplines: computer science, game design, mechanical engineering, physics and psychology. All participants had experience writing and publishing a literature review or a related work section of a research article.

Setup. All interviews were conducted in English via video calls.

Procedure. Each interview lasted 45-60 minutes. After collecting basic information about each scientist’s background, we conducted story interviews, following the same procedure described in Study 1. Participants were asked to describe a recent, concrete example in which they searched for research articles for inclusion in either a literature review or the related work section of an article they were writing.

Data Collection. All interviews were screen recorded and transcribed. We also took hand-written notes and collected examples of documents and software screens.

Analysis Method. The same two researchers analyzed the data using the same reflexive thematic analysis approach described in Study 1. We checked the interview transcripts to reach a shared understanding of the deductive themes for the entire set of interviews, and arrived at agreement after two iterations.

4.2 Results

We collected a total of 33 examples of specific literature search problems encountered by the 12 scientists, approximately three stories per scientist. Although the overall process of searching for the research literature is more open-ended than that of patent examiners, it is similarly complex, and scientists each develop personal strategies for managing their diverse sets of software tools (RQ1). They also develop diverse strategies for capturing and maintaining provenance (RQ2).

4.2.1 Software tool use. One theme related to RQ1 replicates findings from Study 1: *Transferring information across applications is cumbersome (T1)*. We also identified an additional theme related to RQ1: *Reusing notes is difficult (T7)*.

T1: Transferring information across applications is cumbersome. Scientists use specialized applications for reading, note-taking, document management and writing, and many expressed frustration when trying to make these systems work together. For example, every time P5 decides to read a paper in depth, he creates a new Markdown file in Obsidian³. He must then manually create a link between his notes and the source paper. Since many note-taking applications focus on supporting stand-alone reading, scientists struggle to transfer their notes into their final document. Most copy-paste their notes into a text editor, and must then re-organize them and add citations manually. For example, P10 complained that because Overleaf⁴ does not display his personal notes from JabRef⁵, he must re-read the paper to figure out the exact content. All find it difficult to transfer information effectively across different applications.

T7: Reusing notes is difficult. All scientists (12/12) take structured personal notes as they read, which include both objective information, e.g., paper title, link, and quotations, as well as subjective information, e.g., personal summaries and comments. Although most participants (10/12) take notes that summarize the contribution(s) of the article, two (P5 and P8) also copy direct quotations.

All participants (12/12) use multiple applications for recording their notes. Most participants (10/12) comment on the articles directly, using *PDF Reader*'s annotation feature, or write notes on physically printed copies (P9, P11). They also use a diverse set of additional software tools for recording personal notes, each of which imposes a corresponding file format. Some are open-ended, such as *Microsoft Word* (P1, P2, P4, P7, P10, P11), *TextEdit* (P8, P10), and *Emacs* (P6), or hand-written notes (P9) or hand-drawn tree graphs (P3). Others are designed specifically for taking notes, such as *Obsidian* (P5), *Zettlr* (P8) and *Evernote* (P12). The rest are highly structured as *Microsoft Excel* tables (P2, P3, P5, P12).

The format of the original notes directly affects the writing phase. Some (P2, P3, P5, P12) use tables to structure their notes. This provides an overview of the papers, making it easy to compare them (P12), and allowing them to discover “missing bricks in the literature” (P5). However, some uses of tables are problematic. For example, participants find it tedious to manually copy-paste information from multiple sources into each table, and maintain links to sources by




Figure 2: P2 and P5 create tables to compare research articles.

hand. When writing, they must also manually copy-paste contents from the table. They also lack tools for interconnecting their data. For example, P8 wanted the text editor to automatically suggest relevant notes that he could cite. Fig. 2 shows how P2 uses *Excel* to categorize articles according to specified criteria. When she changed her categorization strategy, she decided to color code articles to reflect it, but found it difficult to find relevant notes as she wrote her paper.

More flexible note-writing applications also led to frustration when participants needed to re-organize their notes or keep track of each article's quotations or provenance. For example, P2, P5 and P8 were frustrated by *Word*'s lack of structure which P5 described as a “sandbox” that made it hard to find notes again.

In summary, reusing notes is difficult. Applications that avoid imposing structure during initial note taking make it difficult to find and reuse those notes later. Yet applications that impose an initial structure are only useful if that structure does not change during the sensemaking process.

4.2.2 Provenance management. Two themes related to RQ2 replicate findings from Study 1: *Re-finding sources is tedious (T3)* and that *Search results are easily lost (T4)*.

T3: Re-finding sources is tedious. Scientists have multiple strategies for finding papers, including keyword search (5/12), references from key papers (4/12) and recommendations from others (5/12). Similar to patent examiners, they often need to go back and re-read the papers. For example, P3 re-reads papers when her notes lack sufficient context and she can no longer understand them. Others (P5, P12) manually add a link to facilitate backtracking, but find the process tedious: “When I need to write, I start from my notes, and often go back to the article to see what is corresponding to what.” (P11)

Unlike patent examiners who track individual passages from within documents, most scientists (10/12) focus on notes about the contributions of the whole document. However, two also take fine-grained notes, recording the page number of specific quotations (P9) and adding section-level comments to facilitate refinding (P6).

T4: Search results are easily lost. Participants use diverse search strategies, and want to keep track of their earlier searches. For example, P2 said that the lack of a history of her previous searches within the ACM Digital Library acts as a disincentive: “what I had searched and what rabbit holes I had gone [down] and back.” (P2). P7 was frustrated that he could not remember where he found a particular paper, and thus could not use it as a “memory helper” to start writing.

4.2.3 Summary. Study 2 replicates several findings from Study 1 with respect to the complexity of the document search process

³Markdown (<https://daringfireball.net/projects/markdown>) is a simple mark-up format for text. Obsidian (<https://obsidian.md>) is a “knowledge base” for Markdown files.

⁴<https://www.overleaf.com> – \LaTeX editor

⁵<https://www.jabref.org> – bibliography management tool

across multiple, non-interconnected software tools (RQ1); and the problems associated with managing the provenance of relevant documents (RQ2). Study 2 also identified an additional theme related to RQ1, *Reusing notes is difficult* (T7).

Together, these two studies demonstrate the challenges faced by two types of knowledge workers, patent examiners and scientists, as they read, search, annotate, organize, write and review text documents across multiple specialized applications, while manually tracking their provenance. Although participants in both studies were able to perform their jobs successfully, most reported frustration with the cumbersome nature of the process and the lack of adequate tools.

5 PASSAGES: CONCEPT AND SCENARIO

Based on these findings, we focused our design on facilitating the transfer of information across applications while tracking its provenance. Rather than creating a new, integrated application that supports the complex web of activities we observed, which would have created yet another information silo, we created a new type of interactive object that can be integrated into existing applications.

5.1 Concept

We introduce *Passages*, which seeks to improve inter-connection across applications while maintaining the provenance of information from different sources. Our design approach is an example of “Generative Theory” [8], where a new design is inspired both by empirical studies and by generative principles grounded in theories of human behavior and cognition. Specifically, we were inspired by the principles of reification and reuse [9] from the theory of instrumental interaction [7] to create the concept of a *passage*.

A *passage* is a snippet of text that includes metadata, such as its source document and location within that document, the time it was created, and user-defined tags and comments. A passage can be detached from its source document and reused in other documents and applications. Passages extend the concept of *textlets* [36], which reify a text selection into a first-class object. However, while each

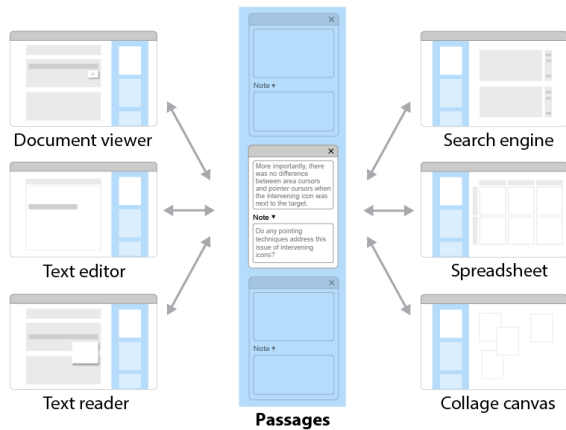


Figure 3: *Passages* reifies text selections into interactive objects that can be manipulated, reused and shared across applications.

textlet is bound to a particular document, passages can be shared **across** multiple applications (Fig. 3). To demonstrate the power of this concept, we created six prototype applications that each support a *Passages Side Panel* for collecting, annotating, manipulating, and sharing passages across applications, without losing their provenance. We first present a scenario to illustrate the user’s experience of using *Passages*, with concrete examples derived the two interview studies. We then describe the system and its implementation in greater detail.

5.2 Scenario

Emma, a Ph.D. student in Human-Computer Interaction, works on improving pointing techniques. Her advisor, Alex, sends her three relevant papers to help her get started: *Silk Cursor* [97], *Area Cursor* [45] and *Enhanced Area Cursor* [94]. Emma must compare them and propose ideas for new interaction techniques. As the process will involve various activities, she decides to use *Passages*.

5.2.1 Reading. She opens the three papers in the Viewer (Fig. 4) and starts to read them, trying to find their similarities and differences. As she reads the *Area Cursor* paper, she realizes that it is inappropriate for fine positioning tasks, since selections may become ambiguous on cluttered displays. She finds this an interesting limitation and collects the corresponding text by selecting it and clicking the **PASSAGE** button. The collected passage immediately appears in the side panel.

As she reads the *Enhanced Area Cursor* paper, she identifies several differences compared to the *Area Cursor*. For example, although the *Enhanced Area Cursor* switches dynamically between pointer and area cursor, performance is the same when an intervening icon appears next to the target. Emma collects this and other relevant text as passages.

5.2.2 Organizing. As the number of passages grows in the side panel, she decides to organize them. She opens the Table (Fig. 4) and drags and drops passages from the side panel into the table. She quickly tires of dragging the passages individually, and decides to move all the passages related to one paper into a column, all at once. She drags the title of the *Enhanced Area Cursor* paper and drops it into an empty column. All the passages automatically fill up the column and she renames the column “Enhanced Area Cursor”.

Emma begins to see patterns in the table. For example, neither *Area Cursor* nor *Enhanced Area Cursor* improve performance when intervening targets appear, and the *Silk Cursor* only works for 3D selection. She summarizes these patterns by naming the rows accordingly. Emma realizes that she has forgotten what one of the passages is about, so she double clicks it, which opens the original source document, with the relevant text selected and highlighted. This helps her examine the passage in its original context. Emma continues to build the table by moving freely among the passages.

5.2.3 Searching. Emma wonders if other papers address the problem of intervening icons next to the target. She returns to the Searcher application (Fig. 4) and revisits her search history, where she discovers a new keyword: “intervening”. She creates a new keyword by selecting the text and clicking the **CREATE** button. She also adds “close” as a synonym and launches a new search with these additional terms. The distribution of the different terms in the

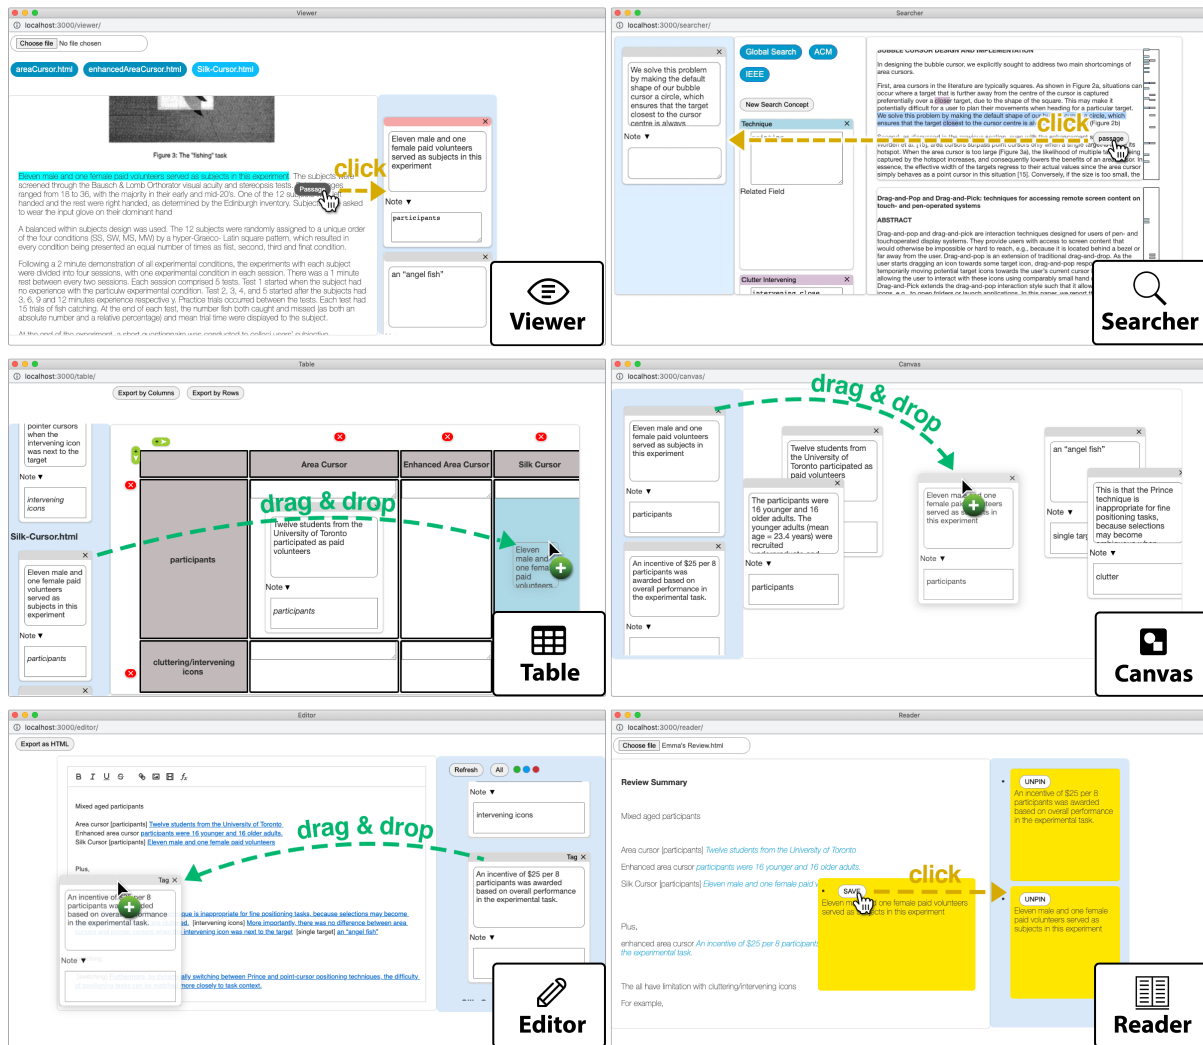


Figure 4: Passages user interface. The **Viewer** lets the user collect interesting text as passages, by selecting them and clicking the **PASSAGE** button. The **Searcher** lets her iteratively search for more documents by specifying multiple search terms, and keeping track of her search history. The **Table** lets her drag and drop collected passages and organize them into rows and columns, whereas the **Canvas** offers a more free-form layout. The **Editor** helps her communicate her findings, while preserving the provenance of each passage as she writes her report. The **Reader** lets others read the report while still easily accessing the source documents to verify the claims.

scrollbar of each resulting document helps Emma quickly locate the relevant text within the paper. After skimming the highlighted sections, she finds the Bubble Cursor paper [33], which seems to address the issue of intervening icons. She creates a new passage with the relevant text description. She then decides to continue reading the document in the Viewer.

5.2.4 Writing. After filling the table with passages, and some additional searching and reading, Emma searches for pointing techniques that remain efficient in the presence of intervening targets. She then decides to write up her findings for Alex. She opens the

Editor (Fig. 4) but does not want to start from a blank page. She returns to the table and uses the *Quick Drafting* feature: she clicks the **EXPORT BY ROW** button (Fig. 5, b), which converts the row header and the passages (together with their notes) from this row into paragraphs in the Editor. Emma re-organizes the text into a concise evidence-based review, with a logical flow. While writing the last paragraph, Emma realizes that one piece of evidence is missing. However, she remembers that she had already read about it in the *Area Cursor Paper*. She opens the paper in the Viewer and immediately finds the related text. She collects it as a passage and drags it from Viewer and drops it into Editor’s side panel (Fig. 5, a). She then inserts it as a citation by dragging and dropping it into the text

area of the Editor. Emma is satisfied with her review and exports it so she can send it to Alex.

5.2.5 Communicating. Alex receives Emma’s review and reads it with the Reader. He wants to find out more about one passage cited as evidence. When he clicks on it, the original passage immediately appears as a yellow tooltip next to the text he is reading. He can also see the full context by double clicking on the passage tooltip, which opens the original paper at the correct scrolling position, with the relevant passage selected and highlighted. This helps him understand Emma’s review. He pins two interesting passages into the side panel, as a reminder that he should discuss them further with Emma at their next meeting.

6 PROOF-OF-CONCEPT IMPLEMENTATION

In order to demonstrate the concept of passage, we created a proof-of-concept implementation with six applications that can share passages: Searcher, Viewer, Table, Canvas, Editor, and Reader.

The *Passages* prototype is implemented as a series of web applications written in HTML, CSS and Javascript with the Vue.js⁶ framework. The back-end is implemented with Node.js⁷ and a NeDB⁸ database, except for the Searcher application, which uses a PostgreSQL⁹ database for full-text search.

6.1 Passages and Passages Side Panels

Each passage is displayed as a box with the quote from the original document and a note that can be edited by the user. It keeps track of both the source document and its position within the document, and can also be tagged with a color picked in a radial menu.

Users can always locate and re-select the text from the passage’s source document simply by clicking on the passage. Since the text is re-selected instead of simply highlighted, users can immediately use normal copy-paste if needed.

A *Passages* side panel (or side panel for short) displays a collection of passages from one or more documents, automatically organized by document. Each of the applications we created includes a side panel, although passages can also exist independently of a side panel. The side panel provides an easy-to-understand, central location for holding the passages relevant to the user’s current task. In addition to automatically classifying passages by their source documents, users can also filter the passages in a side panel according to their color tags.

6.2 Viewer: Collecting Passages

The Viewer application (Fig. 4) offers a lightweight way to collect passages across multiple documents. To create a passage, the user simply selects the desired text with cursor and clicks on the *PASSAGE* button that pops up next to the cursor. The collected passage appears as a persistent interactive object in the side panel, with a small text area for personal notes. We did not provide a categorization mechanism, such as tags, when creating passages to avoid the cost of defining structure too early in the foraging process [49]. However, passages can be tagged once they are created.

⁶VueJS: <https://vuejs.org>

⁷NodeJS: <https://nodejs.org>

⁸NeDB: <https://dbdb.io/db/nebd>

⁹PostgreSQL: <https://www.postgresql.org>

6.3 Table and Canvas: Organizing Passages

The Table and Canvas applications (Fig. 4) are designed to help users organize passages. The Table uses a grid structure, which some participants find particularly useful for comparison tasks [16, 54], while the Canvas offers a more free-form way of organizing passages, which may be more appropriate for long documents or when dealing with a large number of topics.

The Table lets users drag and drop passages back and forth between the table and the side panel, as well as among table cells. Users can populate a column with all the passages related to a specific source document simply by dragging the document title from the side panel to the column. Alternatively, users can click and hold a cell containing a passage to select it all those below it belonging to the same document, and then drag that set to a different cell.

We intentionally separated the collecting application (Viewer) from the organizing applications (Table and Canvas) for two reasons. First, this separation supports the “two-stage” model of sensemaking, where information is gathered first and then organized [49]. Second, this provides more flexibility and opens the possibility of integrating *Passages* into other applications such as *Mural*¹⁰, *LiquidText*¹¹ or *Muse*¹², through a future API.

6.4 Searcher: Searching with Passages

The Searcher application (Fig. 4) combines the designs of *TileBars* [37] and *SearchLens* [17], which let users create multiple keyword objects [17, 37] and visualize them in the scrollbar of each document to support rapid navigation [37]. Unlike *SearchLens* [17], the keyword object has two fields: an editable search query area and a “related” scrapbook. The search query area gives users complete control over their search queries, allowing them to apply advanced search, such as proximity search and add complex synonyms, as suggested by Studies 1 and 2.

The “related” scrapbook addresses the need, identified in the interview studies, to learn about unfamiliar domains by searching for term definitions. Keyword objects let users quickly collect small pieces of information as passages, and attach them to the keyword using drag-and-drop. Users can capture the results of these informal searches and use them to grow their domain knowledge.

6.5 Editor: Writing with Passages

The Editor application (Fig. 4) lets users drag and drop to embed any collected passage as a citation directly into the report. Dropping a passage object inserts that passage’s note as normal text, with the passage itself appearing as a blue citation. This interaction is based on our observations of knowledge workers who copy-paste their notes into their final report for later reuse. Hovering over the blue passage text scrolls the side panel and highlights the corresponding passage object. Users can edit the text as usual, and the link to the passage in the sidebar is preserved. Future work will include an additional button for selecting different citation reference styles, e.g., ACM or APA.

¹⁰Mural: <https://www.mural.co/>

¹¹LiquidText: <https://www.liquidtext.net/>

¹²Muse: <https://museapp.com/>

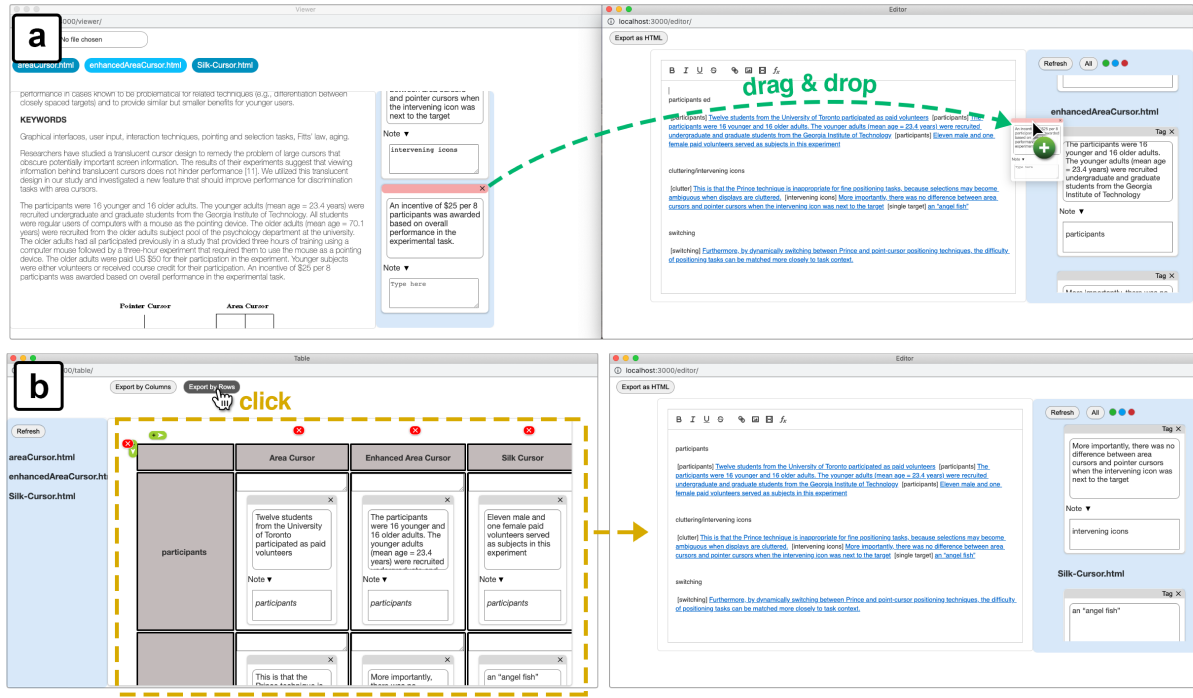


Figure 5: (a) Fluid transitions across applications through drag-and-drop of passages between windows. (b) Reuse of content and structure from the Table directly into the Editor

Table 1: Features of *Passages* that address the research questions (RQ) and themes found in Studies 1 and 2.

RQ	Theme	Passages features
RQ1	T1: Transferring information	Users can drag and drop passages across applications while maintaining their provenance.
RQ1	T2: Recording personal notes	Notes can be attached to passages, and stay attached when moving passages across applications.
RQ2	T3: Re-finding sources	Passages always keep their provenance and let users access the source document with a single click.
RQ2	T4: Searching	Search results can be captured quickly and saved as passages.
	T5: Drafting	Users maintain control when authoring, and can link their writing to passages and their corresponding sources.
	T6: Sharing knowledge	Concept or synonym lists can include and refer to passages, although they cannot currently be shared in real time.
RQ1	T7: Reusing notes	The Table and Canvas provide flexible ways of collecting and organizing notes.

6.6 Reader: Communicating with Passages

The Reader application (Fig. 4) provides a quick and easy way to access any referenced documents. A single click on the blue text lets users see the exact underlying passage as a yellow tooltip. Users can then either open its source document to examine in greater detail, or pin it into the side panel for later reuse. Pinned passages can serve as reminders, e.g. for a discussion with the author. Maintaining a parallel link between the document and its reference information facilitates comparison and verification of the claims, which encourages better evidence-based practices.

We do not include the author’s personal notes in the Reader because patent examiners in Study 1 said that seeing other people’s notes might introduce bias in their own judgement. Future

prototypes should offer the reader the possibility of either seeing or hiding their personal notes.

6.7 Summary

Table 1 summarizes the features of *Passages* that address the themes identified in the Study 1 & 2. Moving passages across applications lets users flexibly interleave knowledge work activities without being forced to follow a preset process. Moreover, the consistent representation of passages across the various applications supports a simple but powerful mental model for users.

Unlike previous systems where note-taking and writing are supported separately, *Passages* lets users reuse their notes as they write while maintaining their provenance. When exporting content from the Table to the Editor, *Passages* automatically collects the row or

column headers and their associated passages, and converts them into text and citations, which users can then edit freely. A more complete implementation should support a similar feature for the Canvas, by letting users specify export rules, e.g., based on tags. Finally the Reader reuses links created during writing process to provide easy access to the original reference material. This connects writing and reading in a fluid cycle, with passages serving as a common underlying objects.

7 STUDY 3: DESIGN WALKTHROUGH WITH PATENT EXAMINERS

We conducted a design walkthrough [59] to gather systematic, focused feedback about the benefits and weaknesses of various features of *Passages*; to collect ideas and suggestions for improvement; and to assess the overall value and relevance of the concept of *Passages*. We based the design walkthrough on a scenario that illustrates the use of *Passages* in a realistic context, with interaction snippets drawn directly from examples provided by patent examiners in Study 1. This scenario is explicitly designed to highlight challenges related to themes identified in Studies 1 and 2.

Participants. We recruited one current patent examiner, two internal software developers, two designers and one manager (two women, four men) from the same patent organization. The two designers were not patent examiners; whereas the other four all had current or previous experience as patent examiners. One developer had participated in Study 1 and was also the author of the automatic drafting application mentioned in Study 1.

Setup. The study was conducted in English via video call.

Procedure. The workshop lasted one hour and featured three main sections: a 5-minute video demo, a 35-minute scenario-based design walkthrough and a 15-minute general discussion. Participants received a description of the design walkthrough procedure prior to the workshop, and a worksheet to fill out with their critiques and suggestions.

After first viewing a 5-minute video of the design scenario described in section 5.2, the first author presented or “walked through” a live demonstration of the scenario, pausing at each step to elicit critiques and suggestions from the participants. This scenario is organized into seven “interaction snippets” that illustrate how a patent examiner uses one or more key features of *Passages* to: 1) collect a passage; 2) organize passages in a table; 3) locate other passages; 4) use the table content for drafting a patent report; 5) include an additional passage; 6) insert a passage as citation; and 7) review the final report. At each step, participants discussed the pros and cons of each feature, and filled out a specially designed worksheet that recorded their critiques and suggestions for improvement. The workshop concluded with a general discussion of the strengths and weaknesses of *Passages*.

Data Collection. We video-recorded the session and collected participants’ worksheets with their critiques and suggestions. We also took hand-written notes during the session.

7.1 Results and Discussion

All participants found *Passages* to be an interesting and elegant concept: “a good way of improving the way we are working now.”(P1). P4 said the concept of interacting with information snippets is

fundamental to their work: “Because we always deal with pieces of information. With this link, you can always navigate between the documents and to communicate etc. I think it is an elegant way of dealing with it. It is not rigid and you can always move it around. This is what I like.” P1 and P3 also appreciated the power and flexibility offered by *Passages*. P1 said: “The difference is really the flexibility, dragging things around, we have no way of doing that.”

Two participants immediately wanted to integrate several features into their organization’s existing applications. For example, P2 wanted to integrate the Reader into their current reviewing applications: “I like the connection that you can open the document. It is really good. I can imagine it working. In terms of integration, actually I could imagine this integrated into the current tools.” P3 especially liked the concept of reusing snippets and the ability to always trace their provenance: “It is a very interesting concept, especially this interaction of having the repository of snippets and just use them in the table and the communication. It saves a lot of time just by linking the documents and having the original version available and having an easy to compare several documents.”

Two participants (P3, P4) liked the connection between reading and writing. P3 liked the ability to quickly access the reference information and the ability to quickly insert a reference in the text editor. P4 said that it is also the ability to have both information (source and reference) in parallel that makes it a great feature. P2 suggested combining reading and writing into a live thread, so that examiners could reply directly within the same passage.

Critiques. Although *Passages* lets users always go back to the source document by double clicking on it, P1 wanted more information about the source document when using a passage in the side panel, Table and Editor, so as to avoid making organization or citation mistakes. This suggests that users should have greater control over configuring the visibility of provenance information. P1 also mentioned that the scalability of the Table may be an issue: “I think I might get lost in a bigger table”. While the current ability to collapse notes may help, we are also considering support for zooming and coloring cells, as in standard spreadsheet software.

Suggestions. Participants viewed the side panel as a “repository of snippets and small fragments” (P3) and found “having them available all the time is really useful” (P6). They suggested adding *bidirectional* connections between the side panel and other applications. For example, P2 suggested that once the passage is included as citation, the side panel could indicate that it has been used. P3 proposed the reverse idea, where users could view all the places in the Editor where a particular passage has been inserted. P2 also suggested connecting the Table and Editor so that passages included as citations in the Editor would appear in the Table as well.

Two participants (P2, P3) immediately understood the power of reusing passages. They brainstormed about structuring mechanisms, such as tagging passages while collecting them. Examiners could then reuse this structure for other tasks, such as filtering and categorizing. P2 also suggested providing Table templates, either predefined or user-defined, with standard clauses in the Editor for examiners to reuse.

Other use cases. Participants suggested other use cases, such as coaching and quality control, where one examiner must review another examiners’ work. P4 also commented on the connection with the literature review process: “I will also generalize it not only

to the review of patent, but also to any academic document, not just limited to patent world.”

8 STUDY 4: STRUCTURED OBSERVATION WITH SCIENTISTS

In order to provide a grounded, qualitative assessment of the strengths and weaknesses of *Passages*, we also conducted a structured observation study [11, 27, 52] with scientists. Participants first perform and then reflect upon a set of realistic literature review tasks, using both their usual literature review process and the *Passages* Viewer, Table and Editor applications. They then compare the details of each, both to identify problems and suggest new ideas. We also asked participants for feedback about *Passages*’s Canvas, Reader and Searcher applications.

Participants. We recruited 12 graduate-level scientists in Human-Computer Interaction (HCI) (9 men, 3 women; ages 22–28; 3 years of research experience on average). They use a diverse set of applications for reviewing literature, including multiple PDF viewers, Google Docs (3/12), Microsoft Word (2/12), Overleaf (2/12), Obsidian (2/12), Google Sheets (1/12), Notion (1/12), and MarginNote (1/12).

Setup. Participants accessed *Passages* via a web application on their personal computers. We prepared two equivalent sets of HCI research papers about basic pointing techniques. Participants were already familiar with *pointing* as an HCI research domain, which is based on a fundamental model—Fitts’ Law—that serves as the “ground truth” for Balakrishnan [5]’s comparative research article.

The first set of three papers is based on the principle of increasing target width: *Area Cursor* [45], *Enhanced Area Cursor* [94] and *Silk Cursor* [97]. The second set of three papers is based on the principle of reducing target distance: *Drag-and-pop* [6], *Pie menus* [13], and *Object pointing* [34]. We use Balakrishnan [5]’s cross-document comparison criteria, specifically participant profile, intervening targets, and commands for switching to normal interaction, to establish the “correct” criteria for analyzing each set of papers.

Procedure. Each session lasts approximately two hours. Participants are asked to read two groups of three research papers, once using their current set of software applications, and once using three *Passages* applications: Viewer, Table, and Editor. They then write a short review that compares the features of pointing techniques for each group of documents.

The two conditions, *document set*: increasing target width or reducing target distance, and *application choice*: with or without *Passages*, are counterbalanced for order within and across participants. Participants hear a scripted presentation that describes the functions of the Viewer, Table and Editor applications immediately before performing the *Passages* condition. We also show participants the Reader application, even though it is not necessary to perform the tasks, so they can see how a fictional colleague would read their report.

The tasks are presented in the form of the following scenario: *Your colleague wants to start a new research project with you. She sends you three documents about pointing techniques that she thinks are relevant but that she has not read. Your goal is to read and compare them. For example, what are the differences and similarities among these techniques? You will write a short paragraph at the end to tell your colleague about what you found out, so she can review*

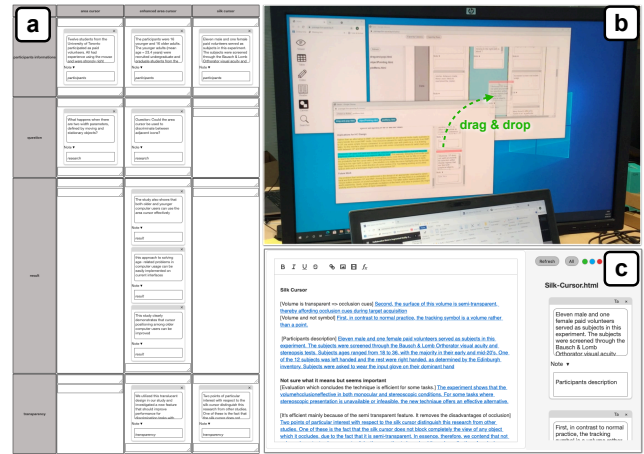


Figure 6: Interacting with *Passages*: Table (a) and report (c) created by participants. (b) P4 drag-dropped a newly created passage directly into the Table.

and verify your findings. In the *Passages* condition, participants perform the task with the Viewer, Table and Editor applications. In the *non-Passages* condition, participants are free to use their own software. Participants use a think-aloud protocol [26] to describe their experiences during each task. After each task, participants complete a short questionnaire and answer questions based on the experimenter’s observations of their behavior.

After completing both tasks, participants complete a final questionnaire that compares their usual literature review process with their experiences using *Passages*. Each question uses a five-point Likert scale, ranging from strongly disagree to strongly agree. Participants are asked to assess their performance on three tasks—making annotations, assessing similarities and differences between documents, and writing an evidence-based report—with respect to their mental load, perceived success, task difficulty and level of frustration. We conclude by asking participants to identify additional use cases from their recent projects; discuss the advantages and disadvantages of using *Passages*; and make suggestions for improvement. We also show participants the Canvas and Searcher applications and ask them to describe situations for which they might be useful.

Data Collection. We collected the passages that participants created with the Viewer, Table and Editor applications, as well as the questionnaire results. We took screen recordings as participants performed each task in both conditions, as well as video of the interviews and discussions. We also took hand-written notes.

8.1 Results and Discussion

We first established the correct criteria for the *increasing target width* condition: participant profile; transparent vs. translucent; and intervening targets; and the *reducing target distance* condition: types of targets; intervening targets; and commands for switching back and forth. We then analyzed the criteria created by each participant to compare the papers in each condition, and counted how many are correct, according to Balakrishnan [5].

All participants successfully completed the task in the *Passages* condition (Fig. 6). Participants considered provenance tracking (8/12) and table exportation (8/12) as their favorite features. Provenance tracking “enables me to work more efficiently in terms of going back to the paper” (P2); “is really good to be able to see the location of citation in the original document” (P4); and “saves me a lot of time finding the context of my previous notes during writing” (P11). P2 liked Table exportation as it “enables me to compare what I have. I like the fact that I can just export it for my writing, not waste it.”

Participants also appreciated the fluid combination of reading and writing: “It feels great because the [Reader is] not something that really exists [in Google Docs]...I really like the way we can click on...and trace the reference” (P1). P3 compared *Passages* to the many note-taking apps she has tried, including *LiquidText*, *MarginNote* and *Notation*, and said none offers *Passages*’ combination of reading and writing.

Participants immediately understood the concept of passages as generic objects that work across multiple applications: “I really like it keeping the notes as separate objects that are not explicitly linked to any PDF, making it a more flexible format (P10).” Some participants wanted to make passages work with their existing applications, such as Mural (P8) and email (P9).

Reuse structure, not just content. Participants reused the passages the collected in the Viewer in the Table (average reuse rate = 77.17%, SD = 0.37) and the Editor (average reuse rate = 88.36%, SD = 0.14). P9 said: “It feels great I can directly use the notes for organization and writing.” Only two participants did not use the table, since they usually write directly into a *Google Doc*.

Participants developed two major strategies for generating their report’s structure: task-based (5/12) and information-based (7/12). Those who used a task-based strategy took advantage of their previous experience with a similar task to generate the initial structure, e.g. pros and cons: “Because the task is to compare interaction techniques, using pros and cons is very natural to me.” (P2). These participants collected passages as they read. By contrast, participants who used an information-based strategy first skimmed two or three papers to gain a general overview, without collecting any passages, since they did not “know what to collect” (P8). They then read each paper in greater detail and began collecting passages.

Despite the differences between these two strategies, all participants wanted to reuse both content and structure. P1 explained: “For me, the important part of using a sheet is the structure.” Similar to the patent examiner’s suggestion in Study 3 to provide tags in the Viewer side panel, both P1 and P6 wanted to tag their passages in the Viewer, and reuse that structure in the Table. Half of the participants (6/12) wanted to reuse structure from the Table in the Editor, either by exporting rows as separate sections or by including column titles as subsections.

Table helps discover new insights. Participants were three times more likely to identify the correct comparison criteria when using *Passages* (Mean = 1.5, SD = 1, total = 18 times) compared to using their own applications (Mean = 0.5, SD = 0.8, total = 6 times). Ten participants found at least one comparison criterion using *Passages*, compared to only four participants who used their own applications. This suggests that *Passages* encourages participants to detect patterns across documents and generate new insights during comparison tasks.

Participants found the Table especially useful in discovering insights and keeping track of their sensemaking process. This includes three participants (P3, P8, P9) who had never previously used tables in their previous literature reviews. P7 said: “The table may help me structure. The format really helps me make something apparent, by adding a table, you have two dimensions.” P6 said the table helped him to have the “epiphany aha moment” about the conceptual link across different domains. P9 said that, while he does not usually work with tables, he was happy he used them because it is very convenient to have all the information at a glance, and described it as a useful “conceptual model”. Three participants who had used their own applications before trying *Passages* spontaneously created a table in that condition using a Markdown editor, *Google sheet* or *LibreOffice*. This supports earlier findings that tables offer a useful representation for analyzing sets of related documents.

Three participants perceived the Table as a temporary placeholder that helped them organize their thoughts. P1 said: “When I’m in a rush, I would completely focus on the writing instead of the sheet.” P3 said she had not previously used tables because she cannot reuse the content when writing, but that the export feature encouraged her to use the Table application.

We also found that several participants (5/12) returned to the Viewer to collect more passages, so they could fill in missing cells in the Table (See Fig. 6, b). P8 said: “The Table feels like a basket and you are looking for eggs.”, and explained that it serves as a description of what he is looking for and guides his search, giving him hints about which particular passages he needs to collect.

Composing and rephrasing. Participants used a writing-oriented strategy when using their own applications, and a composition-oriented strategy when using *Passages*. In the former condition, participants write in their chosen text editor while inspecting their notes or source documents. In the *Passages* condition, participants organise exported text directly into the Editor.

Participants appreciated seeing all their notes in one place, which limits window switching during copy-paste operations and reduces time wasted locating certain notes. P3 said: “In Google Docs, I write mostly by myself. When I’m using *Passages*, I’m less likely to copy and paste contents, they are mostly from the original paper. It’s more convenient. [After importing] I would just change a little, make it more smoothly and logically [organized].”

Participants had mixed opinions about exporting raw passages of text. Most (10/12) said they rephrase the text from the reference document in their own writing to avoid plagiarizing. P10 said: “I’m afraid it will be too similar to the original text if I just copy and paste. You don’t want to plagiarize stuff. Sometimes I quote but I just don’t want to do [it] all the time.” One participant found the exported raw passage a bit “distracting for thoughts and writing something coherent”. Another participant wrote that pie menus have issues with large numbers of items even though the original passage states that “Pie menus seem promising, but more experiments are needed before issuing a strong recommendation.” As he built his table, he realized his mistake and said: “facing the evidence forces me not to twist the words and misquote.” This suggests that passages that appear in the Editor should be collapsible, to reduce clutter while still maintaining access to the original evidence.

Feedback about the Searcher and Canvas applications. Participants quickly saw the benefits of the Searcher for their own work.

For example, P6 had a problem retrieving definitions he had collected about machine learning in his papers, and found it useful to attach definitions to keywords. When P5 was searching for papers related to “immersive technology” and “stereoscopic”, he wanted to use the Searcher to combine these two concepts into a single search. When shown the Canvas, three participants suggested to using it to manage themes in the thematic analysis of their qualitative studies.

Quantitative Results and Limitations. The post-hoc questionnaire asked participants to compare *Passages* to their usual approach, with respect to which is easier, less mentally demanding, less frustrating, and more successful for accomplishing literature review tasks. Although eight participants prefer *Passages* when writing evidence-based results, four found it more frustrating. The follow-up interviews revealed this was because these participants found the Editor’s functionality too limited when compared to a commercial text editor, because it lacks, for example, rich formatting and automatic list formatting.

Participants also suggested adding features to improve the Editor’s citation tracking, such as inserting the document name and maintaining the link after copy-pasting text. Participants considered auto-filling columns in the Table and tagging in the Editor to be the least useful features, although several participants suggested adding tagging to the Viewer, rather than the Editor, since they already have a structure in mind when writing.

Summary. Study 4 demonstrated that scientists can quickly learn to use *Passages* and take advantage of its flexibility. They largely reused the collected passages in terms of both content and structure and used the Table to discover more insights across documents. They also spontaneously devised compositing strategies by organizing exported passages into a concise writing and suggested other use cases, such as thematic analysis.

9 CONCLUSION AND FUTURE WORK

Knowledge work requires interacting with multiple documents to analyze and synthesize information. This paper considers two primary research questions with respect to knowledge work: how do today’s knowledge workers currently perform active reading in their document management process, with particular emphasis on their use of existing software tools (RQ1); and how do they manage the provenance of the documents they use (RQ2)? To address these questions, we interviewed two groups of “extreme” knowledge workers, patent examiners and scientists, and analyzed the complex interconnections among their document-related activities, especially active reading, searching, collecting, organizing, writing and reviewing. We identified seven key difficulties they face, shown in Table 1, five of which relate directly to the above research questions, as well as two that emerged during the thematic analysis of the data. We found that these knowledge workers find it difficult to successfully manage their overall workflow given the set of specialized applications they use (RQ1); and have trouble maintaining the provenance of the information they collect (RQ2).

We then introduce the concept of a *passage*, based on the theory of *instrumental interaction* [7, 9] and its principles of *reification* and *reuse*. A passage is an interactive object that reifies a selected snippet of text and maintains a link to its source document, allowing it to retain its provenance. Each passage can then be manipulated, reused,

and shared across multiple applications using simple interactions such as drag-and-drop. We created a proof-of-concept collection of six applications, each with a *Passages* side panel that supports the creation and sharing of passages across them.

Passages’ ability to support fluid transitions across diverse knowledge work activities, across multiple applications, while maintaining provenance, addresses the above research questions: interconnecting applications rather than creating a new tool (RQ1), and automatically tracking provenance (RQ2). We conducted two studies to assess *Passages*: a design walkthrough with six participants from a major patent organization; and a structured observation study with 12 scientists. Participants found *Passages* to be both elegant and powerful, and especially appreciated its flexibility and their ability to track the provenance of individual passages. They felt it would improve their work practices, facilitate reuse, and offer them novel strategies for analyzing and composing documents. They envisioned new ways of reusing passages for other knowledge work activities, and proposed new ideas, such as bidirectional connections between *Passages* and other applications.

In future work, we plan to extend the current implementation of *Passages* as a stand-alone side panel, so that it can work with the APIs of existing commercial and open-source systems. We also plan to use a platform such as *Webstrates* [50] to support the shareability and distributability of passages across multiple users and devices, so that *Passages* is not limited to static source documents. The *Passages* concept can easily be applied to other text-focused applications, such as community feedback [43, 87]. We also plan to explore how *Passages* can be expanded to handle non-text content, such as transforming selected video clips into passages that can be shared and still maintain their provenance, to support applications such as *VidCrit* [67].

Passages illustrates the power of the deceptively simple operation of “reifying” [7, 9] ephemeral text selections into first class objects that persist and can be transferred across documents and applications without losing their provenance. *Passages* offers a simple yet powerful way to support the ever-increasing complexity of today’s knowledge work, offering flexibility and a fluid iterative workflow.

ACKNOWLEDGMENTS

Many thanks to Domenico Golzio and Jean-Marc Janyszek from the European Patent Office for introducing us to the patent world and for their generous support in establishing the collaboration. We thank all the participants for their rich insights and fruitful discussions. This work was partially supported by European Research Council (ERC) grant № 695464 “ONE: Unified Principles of Interaction”.

REFERENCES

- [1] David Abrams, Ron Baecker, and Mark Chignell. 1998. Information Archiving with Bookmarks: Personal Web Space Construction and Organization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Los Angeles, California, USA) (CHI ’98). ACM Press/Addison-Wesley Publishing Co., USA, 41–48. <https://doi.org/10.1145/274644.274651>
- [2] Eytan Adar, Jaime Teevan, and Susan T. Dumais. 2009. *Resonance on the Web: Web Dynamics and Revisitation Patterns*. Association for Computing Machinery, New York, NY, USA, 1381–1390. <https://doi.org/10.1145/1518701.1518909>
- [3] Annette Adler, Anuj Gujar, Beverly L. Harrison, Kenton O’Hara, and Abigail Sellen. 1998. 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* (Los Angeles, California, USA) (CHI '98). ACM Press/Addison-Wesley Publishing Co., USA, 241–248. <https://doi.org/10.1145/274644.274679>
- [4] Jason Alexander, Andy Cockburn, Stephen Fitchett, Carl Gutwin, and Saul Greenberg. 2009. *Revisiting Read Wear: Analysis, Design, and Evaluation of a Footprints Scrollbar*. Association for Computing Machinery, New York, NY, USA, 1665–1674. <https://doi.org/10.1145/1518701.1518957>
- [5] Ravin Balakrishnan. 2004. "Beating" Fitts' Law: Virtual Enhancements for Pointing Facilitation. *Int. J. Hum.-Comput. Stud.* 61, 6 (Dec. 2004), 857–874. <https://doi.org/10.1016/j.ijhcs.2004.09.002>
- [6] Patrick Baudisch, Ed Cutrell, Daniel Robbins, Mary Czerwinski, Peter Tandler, Benjamin Bederson, and Alex Zierlinger. 2003. Drag-and-Pop and Drag-and-Pick: Techniques for Accessing Remote Screen Content on Touch- and Pen-Operated Systems. In *Human-Computer Interaction-INTERACT '03* (human-computer interaction-interact '03 ed.). IOS Press, 57–64. <https://www.microsoft.com/en-us/research/publication/drag-and-pop-and-drag-and-drag-and-pick-techniques-for-accessing-remote-screen-content-on-touch-and-pen-operated-systems/>
- [7] Michel Beaudouin-Lafon. 2000. Instrumental Interaction: An Interaction Model for Designing post-WIMP User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (The Hague, The Netherlands) (CHI '00). ACM, New York, NY, USA, 446–453. <https://doi.org/10.1145/332040.332473>
- [8] Michel Beaudouin-Lafon, Susanne Bødker, and Wendy E. Mackay. 2021. Generative Theories of Interaction. *ACM Transactions on Computer-Human Interaction* 28, 6, Article 45 (nov 2021), 54 pages. <https://doi.org/10.1145/3468505>
- [9] Michel Beaudouin-Lafon and Wendy E. Mackay. 2000. Reification, Polymorphism and Reuse: Three Principles for Designing Visual Interfaces. In *Proceedings of the Working Conference on Advanced Visual Interfaces* (Palermo, Italy) (AVI '00). ACM, New York, NY, USA, 102–109. <https://doi.org/10.1145/345513.345267>
- [10] Dorrit Billman and Eric A. Bier. 2007. *Medical Sensemaking with Entity Workspace*. Association for Computing Machinery, New York, NY, USA, 229–232. <https://doi.org/10.1145/1240624.1240662>
- [11] Adrien Bousseau, Theophanis Tsandilas, Lora Oehlberg, and Wendy E. Mackay. 2016. How Novices Sketch and Prototype Hand-Fabricated Objects. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 397–408. <https://doi.org/10.1145/2858036.2858159>
- [12] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health* 11, 4 (2019), 589–597. <https://doi.org/10.1080/2159676X.2019.1628806> arXiv:https://doi.org/10.1080/2159676X.2019.1628806
- [13] J. Callahan, D. Hopkins, M. Weiser, and B. Shneiderman. 1988. An Empirical Comparison of Pie vs. Linear Menus. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Washington, D.C., USA) (CHI '88). Association for Computing Machinery, New York, NY, USA, 95–100. <https://doi.org/10.1145/57167.57182>
- [14] Robert Capra, Gary Marchionini, Javier Velasco-Martin, and Katrina Muller. 2010. *Tools-at-Hand and Learning in Multi-Session, Collaborative Search*. Association for Computing Machinery, New York, NY, USA, 951–960. <https://doi.org/10.1145/1753326.1753468>
- [15] Stuart K. Card, George G. Robertson, and William York. 1996. *The WebBook and the Web Forager: An Information Workspace for the World-Wide Web*. Association for Computing Machinery, New York, NY, USA, 111–ff. <https://doi.org/10.1145/238386.238446>
- [16] Joseph Chee Chang, Nathan Hahn, and Aniket Kittur. 2020. Mesh: Scaffolding Comparison Tables for Online Decision Making. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (UIST '20). Association for Computing Machinery, New York, NY, USA, 391–405. <https://doi.org/10.1145/3379337.3415865>
- [17] Joseph Chee Chang, Nathan Hahn, Adam Perer, and Aniket Kittur. 2019. SearchLens: Composing and Capturing Complex User Interests for Exploratory Search. In *Proceedings of the 24th International Conference on Intelligent User Interfaces* (Marina del Rey, California) (IUI '19). Association for Computing Machinery, New York, NY, USA, 498–509. <https://doi.org/10.1145/3301275.3302321>
- [18] James Cheney, Stephen Chong, Nate Foster, Margo Seltzer, and Stijn Vansummeren. 2009. Provenance: A Future History. In *Proceedings of the 24th ACM SIGPLAN Conference Companion on Object Oriented Programming Systems Languages and Applications* (Orlando, Florida, USA) (OOPSLA '09). Association for Computing Machinery, New York, NY, USA, 957–964. <https://doi.org/10.1145/1639950.1640064>
- [19] Harris M Cooper. 1998. *Synthesizing research: A guide for literature reviews*. Vol. 2. Sage.
- [20] Mira Dontcheva, Steven M. Drucker, Geraldine Wade, David Salesin, and Michael F. Cohen. 2006. Summarizing Personal Web Browsing Sessions. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology* (Montreux, Switzerland) (UIST '06). Association for Computing Machinery, New York, NY, USA, 115–124. <https://doi.org/10.1145/1166253.1166273>
- [21] Patrick Dubroy and Ravin Balakrishnan. 2010. *A Study of Tabbed Browsing among Mozilla Firefox Users*. Association for Computing Machinery, New York, NY, USA, 673–682. <https://doi.org/10.1145/1753326.1753426>
- [22] Nathan Evans, Darren Edge, Jonathan Larson, and Christopher White. 2020. News Provenance: Revealing News Text Reuse at Web-Scale in an Augmented News Search Experience. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3334480.3375225>
- [23] Paolo Federico, Florian Heimerl, Steffen Koch, and Silvia Miksch. 2017. A Survey on Visual Approaches for Analyzing Scientific Literature and Patents. *IEEE Transactions on Visualization and Computer Graphics* 23, 9 (2017), 2179–2198. <https://doi.org/10.1109/TVCG.2016.2610422>
- [24] Kristie Fisher, Scott Counts, and Aniket Kittur. 2012. *Distributed Sensemaking: Improving Sensemaking by Leveraging the Efforts of Previous Users*. Association for Computing Machinery, New York, NY, USA, 247–256. <https://doi.org/10.1145/2207676.2207711>
- [25] Martin Flintham, Christian Karner, Khaled Bachour, Helen Creswick, Neha Gupta, and Stuart Moran. 2018. *Falling for Fake News: Investigating the Consumption of News via Social Media*. Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/3173574.3173950>
- [26] Marsha E. Fonteyn, Benjamin Kuipers, and Susan J. Grobe. 1993. A Description of Think Aloud Method and Protocol Analysis. *Qualitative Health Research* 3, 4 (1993), 430–441. <https://doi.org/10.1177/104973239300300403>
- [27] Jérémie Garcia, Theophanis Tsandilas, Carlos Agon, and Wendy E. Mackay. 2014. Structured Observation with Polyphony: A Multifaceted Tool for Studying Music Composition. In *Proceedings of the 2014 Conference on Designing Interactive Systems* (Vancouver, BC, Canada) (DIS '14). Association for Computing Machinery, New York, NY, USA, 199–208. <https://doi.org/10.1145/2598510.2598512>
- [28] Soroush Ghorashi and Carlos Jensen. 2012. Lyleline: Provenance-Based Search Using a Graphical Sketchpad. In *Proceedings of the Symposium on Human-Computer Interaction and Information Retrieval* (Cambridge, California, USA) (HCIR '12). Association for Computing Machinery, New York, NY, USA, Article 2, 10 pages. <https://doi.org/10.1145/2391224.2391226>
- [29] Gene Golovchinsky. 2008. Reading in the Office. In *Proceedings of the 2008 ACM Workshop on Research Advances in Large Digital Book Repositories* (Napa Valley, California, USA) (BooksOnline '08). Association for Computing Machinery, New York, NY, USA, 21–24. <https://doi.org/10.1145/1458412.1458420>
- [30] Gene Golovchinsky, Morgan N. Price, and Bill N. Schilit. 1999. From Reading to Retrieval: Freeform Ink Annotations as Queries. In *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Berkeley, California, USA) (SIGIR '99). Association for Computing Machinery, New York, NY, USA, 19–25. <https://doi.org/10.1145/312624.312637>
- [31] Julien Gori, Han L. Han, and Michel Beaudouin-Lafon. 2020. FileWeaver: Flexible File Management with Automatic Dependency Tracking. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (UIST '20). Association for Computing Machinery, New York, NY, USA, 22–34. <https://doi.org/10.1145/3379337.3415830>
- [32] David Gotz. 2007. The ScratchPad: Sensemaking Support for the Web. In *Proceedings of the 16th International Conference on World Wide Web* (Banff, Alberta, Canada) (WWW '07). Association for Computing Machinery, New York, NY, USA, 1329–1330. <https://doi.org/10.1145/1242572.1242834>
- [33] Tovi Grossman and Ravin Balakrishnan. 2005. *The Bubble Cursor: Enhancing Target Acquisition by Dynamic Resizing of the Cursor's Activation Area*. Association for Computing Machinery, New York, NY, USA, 281–290. <https://doi.org/10.1145/1054972.1055012>
- [34] Yves Guirard, Renaud Blanch, and Michel Beaudouin-Lafon. 2004. Object Pointing: A Complement to Bitmap Pointing in GUIs. In *Proceedings of Graphics Interface 2004* (London, Ontario, Canada) (GI '04). Canadian Human-Computer Communications Society, Waterloo, CAN, 9–16.
- [35] Frank G. Halasz, Thomas P. Moran, and Randall H. Trigg. 1986. Notecards in a Nutshell. *SIGCHI Bull.* 17, SI (May 1986), 45–52. <https://doi.org/10.1145/30851.30859>
- [36] Han L. Han, Miguel A. Renom, Wendy E. Mackay, and Michel Beaudouin-Lafon. 2020. *Textlets: Supporting Constraints and Consistency in Text Documents*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376804>
- [37] Marti A. Hearst. 1995. TileBars: Visualization of Term Distribution Information in Full Text Information Access. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '95). ACM Press/Addison-Wesley Publishing Co., USA, 59–66. <https://doi.org/10.1145/223904.223912>
- [38] Marti A. Hearst. 2009. . Cambridge University Press.
- [39] Morten Hertzum and Annelise Mark Pejtersen. 2000. The Information-Seeking Practices of Engineers: Searching for Documents as Well as for People. *Inf. Process. Manage.* 36, 5 (Sept. 2000), 761–778. [https://doi.org/10.1016/S0306-4573\(00\)00011-X](https://doi.org/10.1016/S0306-4573(00)00011-X)

- [40] Ken Hinckley, Xiaojun Bi, Michel Pahud, and Bill Buxton. 2012. *Informal Information Gathering Techniques for Active Reading*. Association for Computing Machinery, New York, NY, USA, 1893–1896. <https://doi.org/10.1145/2207676.2208327>
- [41] Ken Hinckley, Shengdong Zhao, Raman Sarin, Patrick Baudisch, Edward Cutrell, Michael Shilman, and Desney Tan. 2007. *InkSeine: <i>In Situ</i> Search for Active Note Taking*. Association for Computing Machinery, New York, NY, USA, 251–260. <https://doi.org/10.1145/1240624.1240666>
- [42] Lichan Hong, Ed H. Chi, Raluca Budi, Peter Piroli, and Les Nelson. 2008. SparTag: A Low Cost Tagging System for Foraging of Web Content. In *Proceedings of the Working Conference on Advanced Visual Interfaces* (Napoli, Italy) (AVI '08). Association for Computing Machinery, New York, NY, USA, 65–72. <https://doi.org/10.1145/1385569.1385582>
- [43] Mahmood Jasim, Pooya Khaloo, Somin Wadhwa, Amy X. Zhang, Ali Sarvghad, and Narges Mahyar. 2021. CommunityClick: Capturing and Reporting Community Feedback from Town Halls to Improve Inclusivity. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW3, Article 213 (jan 2021), 32 pages. <https://doi.org/10.1145/3432912>
- [44] Carlos Jensen, Heather Lonsdale, Eleanor Wynn, Jill Cao, Michael Slater, and Thomas G. Dietterich. 2010. *The Life and Times of Files and Information: A Study of Desktop Provenance*. Association for Computing Machinery, New York, NY, USA, 767–776. <https://doi.org/10.1145/1753326.1753439>
- [45] Paul Kabbash and William A. S. Buxton. 1995. The “Prince” Technique: Fitts’ Law and Selection Using Area Cursors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '95). ACM Press/Addison-Wesley Publishing Co., USA, 273–279. <https://doi.org/10.1145/223904.223939>
- [46] David R. Karger and William Jones. 2006. Data Unification in Personal Information Management. *Commun. ACM* 49, 1 (jan 2006), 77–82. <https://doi.org/10.1145/1107458.1107496>
- [47] Amy K. Karlson, Greg Smith, and Bongshin Lee. 2011. *Which Version is This? Improving the Desktop Experience within a Copy-Aware Computing Ecosystem*. Association for Computing Machinery, New York, NY, USA, 2669–2678. <https://doi.org/10.1145/1978942.1979334>
- [48] Alison Kidd. 1994. The Marks Are on the Knowledge Worker. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, Massachusetts, USA) (CHI '94). Association for Computing Machinery, New York, NY, USA, 186–191. <https://doi.org/10.1145/191666.191740>
- [49] Aniket Kittur, Andrew M. Peters, Abdigani Diriye, Trupti Telang, and Michael R. Bove. 2013. *Costs and Benefits of Structured Information Foraging*. Association for Computing Machinery, New York, NY, USA, 2989–2998. <https://doi.org/10.1145/2470654.2481415>
- [50] Clemens N. Klokmoose, James R. Eagan, Siemen Baader, Wendy Mackay, and Michel Beaudouin-Lafon. 2015. <i>Webstrates</i>: Shareable Dynamic Media. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology* (Charlotte, NC, USA) (UIST '15). Association for Computing Machinery, New York, NY, USA, 280–290. <https://doi.org/10.1145/2807442.2807446>
- [51] Janin Koch, Nicolas Taffin, Michel Beaudouin-Lafon, Markku Laine, Andrés Lucero, and Wendy E. Mackay. 2020. ImageSense: An Intelligent Collaborative Ideation Tool to Support Diverse Human-Computer Partnerships. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1, Article 045 (May 2020), 27 pages. <https://doi.org/10.1145/3392850>
- [52] Janin Koch, Nicolas Taffin, Andrés Lucero, and Wendy E. Mackay. 2020. SemanticCollage: Enriching Digital Mood Board Design with Semantic Labels. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference* (Eindhoven, Netherlands) (DIS '20). Association for Computing Machinery, New York, NY, USA, 407–418. <https://doi.org/10.1145/3357236.3395494>
- [53] Siân E. Lindley, Gavin Smyth, Robert Corish, Anastasia Loukianov, Michael Golembewski, Ewa A. Luger, and Abigail Sellen. 2018. *Exploring New Metaphors for a Networked World through the File Biography*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3173574.3173692>
- [54] Michael Xieyang Liu, Jane Hsieh, Nathan Hahn, Angelina Zhou, Emily Deng, Shaun Burley, Cynthia Taylor, Aniket Kittur, and Brad A. Myers. 2019. Unakite: Scaffolding Developers’ Decision-Making Using the Web. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 67–80. <https://doi.org/10.1145/3332165.3347908>
- [55] Bonnie Ma Kay and Carolyn Watters. 2008. Exploring Multi-Session Web Tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy) (CHI '08). Association for Computing Machinery, New York, NY, USA, 1187–1196. <https://doi.org/10.1145/1357054.1357243>
- [56] Wendy Mackay. 2019. Designing with Sticky Notes. In *Sticky Creativity: Post-It Note Cognition, Interaction and Digitalization*, Halskov Christensen and Clemens Klokmoose (Eds.). Academic Press.
- [57] Wendy E. Mackay. 1990. Patterns of Sharing Customizable Software. In *Proceedings of the 1990 ACM Conference on Computer-Supported Cooperative Work* (Los Angeles, California, USA) (CSCW '90). Association for Computing Machinery, New York, NY, USA, 209–221. <https://doi.org/10.1145/99332.99356>
- [58] Wendy E. Mackay. 2000. Responding to cognitive overload: Co-adaptation between users and technology. *Intellectica* 30 (07 2000). <https://doi.org/10.3406/intel.2000.1597>
- [59] Wendy E. Mackay. 2019. Designing with Sticky Notes. In *Sticky Creativity: Post-It Note Cognition, Interaction and Digitalization*, Bo Christensen, Kim Halskov, and Clemens Klokmoose (Eds.). Academic Press, 231–256. <https://hal.archives-ouvertes.fr/hal-02426827>
- [60] Jock D. Mackinlay, Ramana Rao, and Stuart K. Card. 1995. An Organic User Interface for Searching Citation Links. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '95). ACM Press/Addison-Wesley Publishing Co., USA, 67–73. <https://doi.org/10.1145/223904.223913>
- [61] Catherine C. Marshall and Sara Bly. 2005. *Saving and Using Encountered Information: Implications for Electronic Periodicals*. Association for Computing Machinery, New York, NY, USA, 111–120. <https://doi.org/10.1145/1054972.1054989>
- [62] Catherine C. Marshall, Frank G. Halasz, Russell A. Rogers, and William C. Janssen. 1991. Aquanet: A Hypertext Tool to Hold Your Knowledge in Place. In *Proceedings of the Third Annual ACM Conference on Hypertext* (San Antonio, Texas, USA) (HYPERTEXT '91). Association for Computing Machinery, New York, NY, USA, 261–275. <https://doi.org/10.1145/122974.123000>
- [63] Catherine C. Marshall, Morgan N. Price, Gene Golovchinsky, and Bill N. Schilit. 2001. Designing E-Books for Legal Research. In *Proceedings of the 1st ACM/IEEE-CS Joint Conference on Digital Libraries* (Roanoke, Virginia, USA) (JCDL '01). Association for Computing Machinery, New York, NY, USA, 41–48. <https://doi.org/10.1145/379437.379445>
- [64] Kenton O’Hara, Fiona Smith, William Newman, and Abigail Sellen. 1998. Student Readers’ Use of Library Documents: Implications for Library Technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Los Angeles, California, USA) (CHI '98). ACM Press/Addison-Wesley Publishing Co., USA, 233–240. <https://doi.org/10.1145/274644.274678>
- [65] Gerard Oleksik, Natasa Milic-Frayling, and Rachel Jones. 2012. Beyond Data Sharing: Artifact Ecology of a Collaborative Nanophotonics Research Centre. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work* (Seattle, Washington, USA) (CSCW '12). Association for Computing Machinery, New York, NY, USA, 1165–1174. <https://doi.org/10.1145/2145204.2145376>
- [66] Sharoda A. Paul and Meredith Ringel Morris. 2009. *CoSense: Enhancing Sensemaking for Collaborative Web Search*. Association for Computing Machinery, New York, NY, USA, 1771–1780. <https://doi.org/10.1145/1518701.1518974>
- [67] Amy Pavel, Dan B. Goldman, Björn Hartmann, and Maneesh Agrawala. 2016. VidCrit: Video-Based Asynchronous Video Review. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (Tokyo, Japan) (UIST '16). Association for Computing Machinery, New York, NY, USA, 517–528. <https://doi.org/10.1145/2984511.2984552>
- [68] Peter Piroli and Ramana Rao. 1996. Table Lens as a Tool for Making Sense of Data. In *Proceedings of the Workshop on Advanced Visual Interfaces* (Gubbio, Italy) (AVI '96). Association for Computing Machinery, New York, NY, USA, 67–80. <https://doi.org/10.1145/948449.948460>
- [69] Beatriz Pérez, Julio Rubio, and Carlos Sáenz-Adán. 2018. A systematic review of provenance systems. *Knowledge and Information Systems* 57 (12 2018). <https://doi.org/10.1007/s10115-018-1164-3>
- [70] Ramana Rao, Jan O. Pedersen, Marti A. Hearst, Jock D. Mackinlay, Stuart K. Card, Larry Masinter, Per-Kristian Halvorsen, and George C. Robertson. 1995. Rich Interaction in the Digital Library. *Commun. ACM* 38, 4 (April 1995), 29–39. <https://doi.org/10.1145/205323.205326>
- [71] Pamela Ravasio, Sissel Guttormsen Schär, and Helmut Krueger. 2004. In Pursuit of Desktop Evolution: User Problems and Practices with Modern Desktop Systems. *ACM Trans. Comput.-Hum. Interact.* 11, 2 (jun 2004), 156–180. <https://doi.org/10.1145/1005361.1005363>
- [72] George Robertson, Mary Czerwinski, Kevin Larson, Daniel C. Robbins, David Thiel, and Maarten van Dantzich. 1998. Data Mountain: Using Spatial Memory for Document Management. In *Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, California, USA) (UIST '98). Association for Computing Machinery, New York, NY, USA, 153–162. <https://doi.org/10.1145/288392.288596>
- [73] Hugo Romat, Emmanuel Pietriga, Nathalie Henry-Riche, Ken Hinckley, and Caroline Appert. 2019. SpaceInk: Making Space for In-Context Annotations. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 871–882. <https://doi.org/10.1145/3332165.3347934>
- [74] Daniel M. Russell, Mark J. Stefik, Peter Piroli, and Stuart K. Card. 1993. The Cost Structure of Sensemaking. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems* (Amsterdam, The Netherlands) (CHI '93). Association for Computing Machinery, New York, NY, USA, 269–276. <https://doi.org/10.1145/169059.169209>
- [75] Bill N. Schilit, Gene Golovchinsky, and Morgan N. Price. 1998. Beyond Paper: Supporting Active Reading with Free Form Digital Ink Annotations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Los Angeles, California, USA) (CHI '98). ACM Press/Addison-Wesley Publishing Co., USA, 249–256. <https://doi.org/10.1145/274644.274680>

- [76] M. C. schraefel, Yuxiang Zhu, David Modjeska, Daniel Wigdor, and Shengdong Zhao. 2002. Hunter Gatherer: Interaction Support for the Creation and Management of within-Web-Page Collections. In *Proceedings of the 11th International Conference on World Wide Web* (Honolulu, Hawaii, USA) (WWW '02). Association for Computing Machinery, New York, NY, USA, 172–181. <https://doi.org/10.1145/511446.511469>
- [77] Abigail J. Sellen, Rachel Murphy, and Kate L. Shaw. 2002. How Knowledge Workers Use the Web. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Minneapolis, Minnesota, USA) (CHI '02). Association for Computing Machinery, New York, NY, USA, 227–234. <https://doi.org/10.1145/503376.503418>
- [78] Nikhil Sharma. 2011. *Role of Available and Provided Resources in Sensemaking*. Association for Computing Machinery, New York, NY, USA, 1807–1816. <https://doi.org/10.1145/1978942.1979204>
- [79] Frank M. Shipman, Catherine C. Marshall, and Thomas P. Moran. 1995. Finding and Using Implicit Structure in Human-Organized Spatial Layouts of Information. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '95). ACM Press/Addison-Wesley Publishing Co., USA, 346–353. <https://doi.org/10.1145/223904.223949>
- [80] David Canfield Smith, Charles Irby, Ralph Kimball, and Eric Harslem. 1982. The Star User Interface: An Overview. In *Proceedings of the June 7-10, 1982, National Computer Conference* (Houston, Texas) (AFIPS '82). Association for Computing Machinery, New York, NY, USA, 515–528. <https://doi.org/10.1145/1500774.1500840>
- [81] Craig A. N. Soules and Gregory R. Ganger. 2005. Connections: Using Context to Enhance File Search. In *Proceedings of the Twentieth ACM Symposium on Operating Systems Principles* (Brighton, United Kingdom) (SOSP '05). Association for Computing Machinery, New York, NY, USA, 119–132. <https://doi.org/10.1145/1095810.1095822>
- [82] Kate Starbird, Dharma Dailey, Owla Mohamed, Gina Lee, and Emma S. Spiro. 2018. *Engage Early, Correct More: How Journalists Participate in False Rumors Online during Crisis Events*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3173574.3173679>
- [83] S. Stumpf, E. Fitzhenry, and T. G. Dietterich. 2007. The use of provenance in information retrieval. In *Workshop on Principles of Provenance (PROPR)*. 2. <https://openaccess.city.ac.uk/id/eprint/224/>
- [84] Hariharan Subramonyam, Colleen Seifert, Priti Shah, and Eytan Adar. 2020. *TexSketch: Active Diagramming through Pen-and-Ink Annotations*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376155>
- [85] Craig S. Tashman and W. Keith Edwards. 2011. *Active Reading and Its Discontents: The Situations, Problems and Ideas of Readers*. Association for Computing Machinery, New York, NY, USA, 2927–2936. <https://doi.org/10.1145/1978942.1979376>
- [86] Craig S. Tashman and W. Keith Edwards. 2011. *LiquidText: A Flexible, Multitouch Environment to Support Active Reading*. Association for Computing Machinery, New York, NY, USA, 3285–3294. <https://doi.org/10.1145/1978942.1979430>
- [87] Sunny Tian, Amy X. Zhang, and David Karger. 2021. A System for Interleaving Discussion and Summarization in Online Collaboration. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW3, Article 241 (jan 2021), 27 pages. <https://doi.org/10.1145/3432940>
- [88] E von Hippel. 1986. Lead Users: A Source of Novel Product Concepts. *Manage. Sci.* 32, 7 (jul 1986), 791–805.
- [89] Julie Wagner, Wendy E. Mackay, and Stéphane Huot. 2012. Left-over Windows Cause Window Clutter... But What Causes Left-over Windows?. In *Proceedings of the 2012 Conference on Ergonomie et Interaction Homme-Machine* (Biarritz, France) (*Ergo'12*). Association for Computing Machinery, New York, NY, USA, 213–216. <https://doi.org/10.1145/2652574.2653442>
- [90] Harald Weinreich, Hartmut Obendorf, Eelco Herder, and Matthias Mayer. 2008. Not Quite the Average: An Empirical Study of Web Use. *ACM Trans. Web* 2, 1, Article 5 (March 2008), 31 pages. <https://doi.org/10.1145/1326561.1326566>
- [91] Ryen White and Resa Roth. 2009. . Morgan & Claypool Publishers. <https://doi.org/10.2200/S00174ED1V01Y200901ICR003>
- [92] Steve Whittaker and Candace Sidner. 1996. *Email Overload: Exploring Personal Information Management of Email*. Association for Computing Machinery, New York, NY, USA, 276–283. <https://doi.org/10.1145/238386.238530>
- [93] Allison Woodruff, Andrew Faulring, Ruth Rosenholtz, Julie Morrisson, and Peter Pirolli. 2001. Using Thumbnails to Search the Web. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Seattle, Washington, USA) (CHI '01). Association for Computing Machinery, New York, NY, USA, 198–205. <https://doi.org/10.1145/365024.365098>
- [94] Aileen Worden, Nef Walker, Krishna Bharat, and Scott Hudson. 1997. Making Computers Easier for Older Adults to Use: Area Cursors and Sticky Icons. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (CHI '97). Association for Computing Machinery, New York, NY, USA, 266–271. <https://doi.org/10.1145/258549.258724>
- [95] Dongwook Yoon, Nicholas Chen, and François Guimbretière. 2013. TextTearing: Opening White Space for Digital Ink Annotation. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (St. Andrews, Scotland, United Kingdom) (UIST '13). Association for Computing Machinery, New York, NY, USA, 107–112. <https://doi.org/10.1145/2501988.2502036>
- [96] Dongwook Yoon, Nicholas Chen, François Guimbretière, and Abigail Sellen. 2014. RichReview: Blending Ink, Speech, and Gesture to Support Collaborative Document Review. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (Honolulu, Hawaii, USA) (UIST '14). Association for Computing Machinery, New York, NY, USA, 481–490. <https://doi.org/10.1145/2642918.2647390>
- [97] Shumin Zhai, William Buxton, and Paul Milgram. 1994. The "Silk Cursor": Investigating Transparency for 3D Target Acquisition. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, Massachusetts, USA) (CHI '94). Association for Computing Machinery, New York, NY, USA, 459–464. <https://doi.org/10.1145/191666.191822>
- [98] Amy X. Zhang, Michael Muller, and Dakuo Wang. 2020. How Do Data Science Workers Collaborate? Roles, Workflows, and Tools. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1, Article 022 (May 2020), 23 pages. <https://doi.org/10.1145/3392826>
- [99] Xiaolong Zhang, Yan Qu, C. Lee Giles, and Piyong Song. 2008. CiteSense: Supporting Sensemaking of Research Literature. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy) (CHI '08). Association for Computing Machinery, New York, NY, USA, 677–680. <https://doi.org/10.1145/1357054.1357161>