

IRCHIVER: An Information-Centric Personal Web Archive for Revisiting Past Online Sensemaking Tasks

Zhicheng Huang

Brown University
Providence, Rhode Island, USA
zhicheng_huang@brown.edu

Jeff Huang

Brown University
Providence, Rhode Island, USA
jeff_huang@brown.edu




Fig. 1. IRCHIVER's workflow. (1) It captures browser screenshots along with their corresponding page sources; (2) The index server stores on-screen text (extracted via OCR) as well as off-screen text and metadata (both parsed from page sources) in a local inverted index. It also runs a compressor daemon that converts screenshots from PNG to WEBP; (3) Finally, the search server hosts a local search interface powered by the inverted index.



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

CHIIR '25, Melbourne, VIC, Australia

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1290-6/25/03

<https://doi.org/10.1145/3698204.3716459>

Abstract

The dynamic nature of web content poses unique challenges to revisiting past online sensemaking tasks. We investigate the value and usage patterns of a personal web archive that treats consumed information as first-class object in both information re-finding and mental model reconstruction. In this work, we introduce a new system called IRCHIVER that passively captures screenshots of active browser windows, extracts text from screenshots and page sources,

and indexes consumed web information to be searched later. To evaluate the effectiveness of IRCHIVER, we conducted a mixed-method study where we compare people's perceived values in mental model reconstruction and information re-finding between IRCHIVER and browser-native features. Our user study suggests that people are able to re-find information more effectively, restore mental models more completely, and reconstruct decision-making processes with greater accuracy, and overall revisit past sensemaking tasks more confidently. Additionally, semi-structured interviews with participants illustrate implications for the design of user interface for personal web archives.

CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; • **Information systems** → *World Wide Web*.

Keywords

Personal Information Management, Web Information Re-finding, Mental Model Restoration

ACM Reference Format:

Zhicheng Huang and Jeff Huang. 2025. IRCHIVER: An Information-Centric Personal Web Archive for Revisiting Past Online Sensemaking Tasks. In *2025 ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR '25)*, March 24–28, 2025, Melbourne, VIC, Australia. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3698204.3716459>

1 Introduction

Whether it is journalists skimming headline articles on a news platform to gauge market sentiment, or graduate students comparing the trade-offs of different flight itineraries to plan a trip, people often need to make sense of dynamic web content that is constantly updating. Often times, people need to revisit what they have seen earlier [29, 71, 95, 99]. Revisiting a past sensemaking task usually involves more than information re-finding. In addition to re-finding online information one looked at, it is also important to restore the mental model associated with that information because mental models are an integral part of task resumption [49, 50, 70, 86, 87], collaborative sensemaking [63, 75], and metacognitive processes such as self-reflection [51, 80]. For example, besides re-finding the original decided flight itinerary, one would need to restore the mental model of why that flight was better than other alternatives in order to explain the decision to others.

However, it is difficult to re-find information with the dynamic nature of online information. Pages that are important to one's past sensemaking tasks often undergo changes over time [8, 69, 82, 89, 90, 106]. Dynamic content often leads to inefficient re-finding and a sense of frustration [71, 89, 98]. For example, the original navigation path leading to the search result has been shown to be important for people re-finding the same result [19, 100, 104]. However, when a web page used by the user as a "waypoint" to reach the final web resource undergoes content drift after the initial search, it would be challenging to re-find the information originally saw. Public archiving services such as the Wayback Machine [55] periodically take screenshots of web resources. Although these snapshots provide original information even in the presence of link rots, the coverage of snapshots is still limited [4]. Moreover, they

can only curate the public view of the page and not the stateful or manipulated pages that users see. For example, they cannot store snapshots of the Wall Street Journal because it requires subscription before accessing full contents, and they cannot store snapshots of the list of available flight options on Google Flights because it required states to reach that page, including dates, destination, and various filter options.

Browsers offer ways for users to re-find information. People often have specific information in mind during re-finding, but browser history tools only store page URLs along with a few metadata (e.g., title, domain name, and favicon) and are therefore ineffective [29, 71, 95]. Existing Personal Information Management (PIM) systems tried to enhance users' information re-finding experience by incorporating browsing activities [56, 80], context [30, 32, 35, 37, 101], metadata [46, 83], or active manual curation [12, 64, 66]. However, they were often limited in one or more of the following ways: (1) does not index on-page content for users to search, (2) acts barely as a pointer to the page, which shows the latest content instead of the original content, and (3) does not directly promote mental model restoration beyond relying on users' memory.

Building on prior research demonstrating that visual cues from screenshots can effectively evoke situational context [51, 86, 87], can we design a system that supports both information re-finding and mental model reconstruction, enabling users to dynamically revisit and engage with past sensemaking tasks in a web environment?

In our work, we explore the idea of having a system passively helping users "remember" what they see during browsing and make it easy for them to revisit past sensemaking tasks, including information re-finding and mental model restorations. We instantiate this idea in a system called IRCHIVER, a full-resolution searchable personal web archive of what people see during browsing that treats information as a first-class object. IRCHIVER captures images of pages viewed in the user's web browser through a background process on Windows and macOS computers, automatically converts them into a compressed WEBP format, and extracts text via optical character recognition (OCR) as well as from the page source for indexing and searching. Besides archiving browsing history, IRCHIVER naturally stores stateful representations of websites, including snapshots taken while forms are filled out or after elements on the website have changed. Essentially, everything seen through the user's browser can be saved, searched, and preserved using a modest amount of space. For the user, this creates a "search history" tool, providing an archive of all pages they have seen and a personal search engine to revisit past online tasks. The archiving capabilities of IRCHIVER are fully automated and operate passively in the background.

We conducted a user study to evaluate the effectiveness of IRCHIVER in helping people revisit past sensemaking tasks and to understand how they use such a system. We found that IRCHIVER helped participants re-find information more efficiently, restore mental models more completely, and reconstruct decision-making processes more accurately, thereby increasing their confidence during revisitation.

The primary contributions described in this paper include:

- a cross-platform system called IRCHIVER that passively curates on-screen content into a searchable personal web archive with full-resolution screenshots
- quantitative and qualitative evidence demonstrating the value of visual snapshots and full-text search for revisiting past online sensemaking tasks among dynamic web content
- design implications for future passive personal web archives that promote both information re-finding and mental model reconstruction

2 Related Work

2.1 Re-finding Online Information

Re-finding online information is a common task. Teevan et al. [99] concluded that re-finding queries constituted as many as 40% of all queries. In studies that analyzed users' browser logs and click-stream data, page re-visitation rate was reported to be 44% [71], 58% [95], and 81% [29]. However, it is challenging to re-find online information. Wen [104] found that the overall success rate for study participants to re-find visited pages was below 20%. One reason for the difficulty of re-finding information is that it is a challenge for people to keep track of how information is organized in web repositories such as bookmarks [1, 29].

The dynamic nature of web information makes re-finding online information more difficult. Content drift [8] refers to the phenomena where owners of web resources add, retract, or replace. This includes the updated results for the same query in search engine result pages (SERP) [89]. In the extreme case, the web resource pointed by a URL can be removed, a scenario commonly known as link rot [69, 82, 90, 106]. Previous studies show that the original path that leads to the search result is important for people to re-find the same result [19, 100, 104].

Browsers offer ways for users to re-find information. Browser-native tools such as history and bookmarks store page URLs along with a few metadata (e.g. title, domain name, and favicon). However, Jones et al. [58] observed that browser-native tools, including browser history and bookmarks, were not frequently used by participants. This was consistent with what Obendorf et al. [71] observed in their study where history and bookmarks were used in only 16% of medium and long-term revisitations. Opening pages with new tabs and new windows is a commonly used re-finding strategy. Dubroy and Balakrishnan identified tab switching as the second most frequently used navigation method, behind link clicking, and that a median of 73.3% of tab switches fell under revisitations [34]. Likewise, Chang et al. [22] showed users decided to leave tabs open to avoid re-finding costs. Weinreich et al. [103] suggested that tabbed browsing became popular because it helped users avoid the need to backtrack. However, opening more tabs and windows leads to cluttered browsing [67] and contributes to information overload [40].

2.2 Personal Information Management

Vannevar Bush pioneered the idea of "memex", which is a device that an individual can store all the books, records, and communications for later consultation [15]. Adar et al. [2] introduced the idea of curating a "haystack" of personal information for each individual that supported annotations and collections. Dumais et al. [35] built

a unified index of information on users' personal computers where rich contextual cues such as time, thumbnails of the application, and authors were used to facilitate information re-finding. Compared to the screenshots presented by IRCHIVER, its thumbnails were more similar to the favicons shown in modern browser history tools. In a follow-up study, Cutrell et al. [30] presented a more robust searching interface that supported various filtering and tagging, making it easier for users to organize their personal information. However, changes to objects posed serious problems to the system due to update latency. Gemmell et al. [39] presented a PIM for multimedia files where it offered both a timeline view and a clustered-time view, which cluster images by similar time and arrange them in time order, for query results. Morris et al. [70] structured search topics, queries, results, and users' annotations in a hierarchical structure to facilitate information re-finding. By presenting terminal pages under the query that led to them, it showed the context for each page. Prior research tried to use cached versions of users' information space to "time travel" to earlier information environments [44, 81]. Teevan [97] designed a search engine that incorporates cached search results of a query to the latest fetched results. While this approach mitigates the difficulty content drift poses to re-finding tasks through reconstructing original paths, it is still subject to link rot. To facilitate efficient web archiving, Pehlivan et al. [76] incorporated visual representation of web pages in detecting web page changes.

Sometimes metadata are saved to facilitate search. Since individuals have different ways to browse the internet, personalization offers benefits to quickly locating relevant browsing history and clusters of interested topics. Memex for the Web actively saves users' browsing history, metadata, keywords, and web page structures from synthesizing personalized topic clusters [21]. The system enables users to search coherent topics from their prior history through relevant topics. In addition to metadata, context was also used to help users re-find the original query [32, 74].

2.3 Web Archiving Tools

IRCHIVER can also be compared to other modern archiving methods (Tab. 1). Users can take a screenshot or save the page they are viewing as a PDF, which are manually initiated and store what they see into a local file. They need to organize and index those files manually if they want them to be retrievable later, and the need to manually capture each page means that most pages will not be saved. Some software tools such as WebRecorder [102] and ArchiveBox [5] can capture using software installed locally, but while they may have more access to what was seen by the user, they still require manual intervention for each page. On the other hand, there are cloud-based archivers that crawl and store other pages from their own servers, such as the Wayback Machine [55], archive.org [54], and archive.ph [6]. These services visit the target link independently, and archive what was presented to them. This means that it is possible to do it automatically, which archive.org does, but they are only able to get the public view of the page, and not the stateful or manipulated pages that users see.

Archiving large-scale web content for full-text search lets users engage with web content directly and refine the search scope [11].

Tab. 1. A comparison of IRCHIVER’s capabilities with those of current browser archiving options.

Product	When Captured	Storage	Format	Reproduction	Platform
IRCHIVER	Continuous, automatic	Local	Image, text	What you saw	Cross-platform, any browser
Archive.org	Occasional, automatic	Cloud	Webpage	Only public view	Web
Archive.ph	Manually initiated	Cloud	Webpage	Only public view	Web
ArchiveBox	Manually initiated	Local	Webpage	Only public view	Any desktop
WebRecorder	Manually initiated	Local	Webpage	What you scroll	Chrome
Taking a screenshot	Manually initiated	Local	Image	What you saw	Any desktop
Browser extensions	Manually initiated	Local	Varies	Varies	Varies
Browser “Save as PDF”	Manually initiated	Local	PDF	Nearly what you saw	Most browsers

Ben-David and Huurdeman note that searching web archives allow for “the comparative analysis of search results, the study of the standing of issues or topics in particular points in time, the identification of archival artifacts, or the reassemblage of existing collections based on thematic, temporal or technical criteria.”

2.4 User-Contributed Browser Logging Tools for Researchers

Multiple tools developed by information retrieval researchers have attempted to gather user-contributed browsing logs for researchers. Perhaps the most well known is the Lemur Query Log Toolbar, mentioned in the 2009 TrebleCLEF workshop report [28], but developed as an add-on that anyone could install and contribute anonymized browsing data to researchers. Unfortunately, not enough people contributed to the project, and it ceased to operate without releasing any of the data. CrowdLogger by Feild and Allen was built off it [38], and used in small efforts by the authors for their own studies, but also never reached public use. Capra’s HCI Browser [17] was a browser plug-in for Firefox, mainly meant for administrating web studies, but it is no longer functional. Meanwhile, Pulliza and Shah’s IRIS [79] was a browser plugin for multifaceted logging, with functionality intended for search sessions, but it has been discontinued and it’s not clear if regular users discovered and benefited from this tool. Perhaps closest to this work is WASP by Kiesel et al. [61], a personal web archive and search system. It works by injecting into the browser’s stream to capture the page source as it is transmitted from the online source to the user’s web browser. The technical sophistication makes it difficult for a user to setup, but its capture of page source makes it more versatile. However, Kiesel et al. [62] admits that “the reproduction of web pages from archives is far from perfect” when evaluating their automated archive quality assessor on a corpus of archived pages.

2.5 Mental Reconstruction Promoted by Visual History

Visual cues have been widely used to invoke memories and help users recognize previously visited documents. Kaasten et al. [59] found that a visual thumbnail could provide context to make a page recognizable where textual content helped users pinpoint a particular page. Teevan et al. [101] explored various representations of web pages and found that augmenting a visual thumbnail with a textual snippet offered the most effective search support while remaining significantly smaller than a purely text-based snippet. Organizing thumbnails in spatial [83] and hierarchical [47] leveraged people’s




Fig. 2. The options menu displayed when the user clicks the IRCHIVER icon on their desktop. This menu provides quick access to the search interface and local repository that stores users’ data.

ability to associate content with location and improved user’s ability to re-find saved pages compared to traditional browser bookmarks.

Previous research showed that visual histories of activities are an effective way to help users reconstruct their working contexts. Cangiano et al. [16] found watching screen recordings of past work helped people remember contextual details such as why they were working on a particular project. Recently, Rule et al. [87] and Hu et al. [50] confirmed the benefits of using visual history to restore situational context, in the settings of task resumption. With the context invoked by visual cues, Hu et al. found that visual history could also promote self-reflection and perceived productivity among users [51].

3 System Design

3.1 Background Process

IRCHIVER starts when the computer boots up, and it runs in the background. It remains completely silent but displays an icon in the system notification area (for example, the bottom-right corner on Windows or the menu bar on macOS). This icon indicates that there is background capturing happening, as a cue for users to prevent invisibly capturing images secretly, e.g. if the computer user is unaware that it’s there. The icon also offers some options for accessing the features in IRCHIVER (Fig. 2).

By being a background process, IRCHIVER is more robust than a browser extension, which depends on the permission of the browser




Fig. 3. The search interface of IRCHIVER. (A) Search configuration allows users to include OCR-processed content or the full page source in their queries; (B) Sorting options enable users to sort results by relevance or time; (C) Date filters provide the ability to specify a date range for searches; (D) Full-text search bar for entering search queries; (E) Results panel displays each search result with a screenshot and associated metadata, including the URL, title, and timestamp of the visit.

vendors and the architecture of the browsers’ support for extensions. In comparison, the desktop software is expected to work for a long time, as the Windows API and macOS Accessibility functionality have been supported for a long time with backward compatibility. The native desktop application format is private by default; both the client and the server are integrated into this application file that runs only on the user’s computer (Fig. 1); no data is ever sent outside of the user’s computer without the user explicitly transferring it themselves, as IRCHIVER by default has no networking capabilities. IRCHIVER is publicly available¹.

3.2 Image Capture Mechanism (Client)

The images from the browser are captured by an event hook, where IRCHIVER monitors the browser’s address bar and viewing windows to detect when something has changed. All major browsers are supported by this technique: Chrome, Firefox, Edge, Brave, and Opera. When a change occurs, IRCHIVER determines whether enough of a difference has happened to decide whether to make a new capture; a configurable combination of the amount of changes and the time since the last change, typically leads to each new URL being captured at least once, and multiple times if the user is scrolling on

the page. The captures are made on the image even if rendered by hardware acceleration using an undocumented but long-supported Windows API function, as typical methods of capture will not work on most browsers which render the pages in hardware. The macOS implementation adopts a similar strategy, using event observers in the Accessibility API alongside *ScreenCaptureKit*. The captures are quietly stored on the user’s local desktop as plain text and images, making it easy for users to review the data even without IRCHIVER.

One intentional design decision is whether to capture images, or to capture the page source to later re-render the page. Unlike previous beliefs, we believe that images (i.e. screenshots) are the most reliable representation, as they are rendered by the exact version of that browser, with a particular browser configuration in that moment, and is exactly what was seen by the user in their sized browser window. Page source can be rendered into interactive pages, but ends up being re-rendered in a different environment (e.g. installed fonts, browser window size, state of the running JavaScript code), and may have problems if data needs to be retrieved from elsewhere online, or if the state is not recreated exactly right. The resulting decision was to capture images as the primary form of storing history, which differs from how most web archivers work. Notably, while IRCHIVER does not capture the entire page source

¹<https://irchiver.com/>

for re-rendering, it does collect the page source to complement the screenshots with additional context.

IRCHIVER can therefore capture stateful pages, even after the user interacts with them or input data after they were rendered. While page source is much smaller, WEBP provides a sufficient trade-off between size and compression, where most 4K resolution captures are significantly under 1 MB without any lossy compression, and the largest images with photorealistic content on the page can be compressed down to under 1 MB without any noticeable degradation. Computers are now fast enough to do this compression in real-time without affecting other processes. By clicking the “Show Data Folder” option in IRCHIVER’s menu (Fig. 2), users can inspect the local repository that stores their local data. When users need to free up disk space or are certain they won’t return to specific data, IRCHIVER offers an easy way to inspect and delete it locally.

3.3 Dual Indexing of OCR Text and Page Source (Server)

In a background process running on the user’s computer, the IRCHIVER index server monitors the filesystem for newly captured images and page source files. When such files appear, it indexes them into a SQLite-backed inverted index, scans the images for text using the OCR library *tesseract*, and processes the textual content from the page source concurrently.

One key distinction between the text captured by OCR and that extracted from the page source lies in their respective scopes. OCR processes only the text visible on-screen, potentially missing off-screen content that users have not scrolled into view. In contrast, extracting text from the page source can omit text embedded within images (e.g., infographics or video captions) and dynamically generated content. Consequently, while there is significant overlap between the text obtained via OCR and that parsed from the page source, each method captures distinct aspects of the webpage content.

3.4 Search Interface (Server)

The IRCHIVER search interface is simple (Fig. 3), allowing the user to choose which index (OCR’d text or page source) to search, or even both, and shows thumbnail results alongside the page title, link, and time captured. Clicking on the thumbnail opens the WEBP image in the browser, which looks pixel-for-pixel the same as the original page. Pages from the same link are clustered together, as one page visit may generate multiple images, such as if the user was scrolling on the page, or spent a long time on it writing an email. By default, IRCHIVER uses BM25 [84], a classic lexical search algorithm, to rank the results. It also supports sorting by time, so the results can be presented in a timeline view.

4 Methods

We conducted a user study to evaluate the effectiveness of IRCHIVER in helping people revisit past sensemaking tasks. We aimed to address the following research questions:

RQ1: What values do people get from an information-centric PIM system?

RQ2: How do people revisit complex online sensemaking tasks with an information-centric PIM system both in re-finding information and restoring mental models?

RQ3: How would people integrate an information-centric PIM system into existing workflows?

4.1 Participants

We recruited 14 participants (10 male, 4 female) aged 20–29 ($\mu = 24.7$, $\sigma = 6.2$) years old through email lists, social media, and flyers. The participants were required to be 18 or older, fluent in English, and have a personal computer. Participants spend an average of 7.6 hours using browsers every day, eleven participants perform decision-making tasks (beyond simple information lookups) on a daily basis and the rest on a weekly basis. Two participants used Windows computers and twelve used MacOS laptops. Eight participants used Google Chrome, three used Firefox, two used Edge, and one used Safari in the study.

4.2 Procedure

The study consisted of two 60-minute sessions. All studies were completed either in-person or remotely online over Zoom, where participants were asked to share their screens under both modes of participation. The IRB of the searchers’ institution decided the protocol meets the criteria for exemption from IRB review. Participants were compensated at \$20 per hour for all user studies.

4.2.1 Overview. Fig. 4 shows an overview of the study procedure. The study employed a within-subject design, with two sessions conducted 2 to 7 days apart, based on participants’ availability. In the first session, participants completed four sensemaking tasks across two scenarios, each containing two tasks, in which they consumed online information and made decisions based on their findings. In the second session, participants were asked to revisit each of the initial four sensemaking tasks including restoring original mental models and re-finding pages where they learned relevant factors that contributed to their original decisions. In the second session, for the two tasks under each scenario, participants were asked to (1) use IRCHIVER to revisit the initial sensemaking task and (2) any reasonable tool(s) except for IRCHIVER to revisit the other task.

4.2.2 Justification for Time Interval and Task Selection. To emulate sessions that take place a period of time apart, a time interval of 2–7 days was used to separate the two sessions that each participant had. The selection of this interval was informed by both memory theory and practical considerations. The Ebbinghaus forgetting curve indicates that the most significant decline in memory retention occurs within the first two days, stabilizing afterward: people typically retain approximately 28% of learned information after 2 days and about 25% after 7 days [36]. This observation was confirmed by pilot studies, which showed that participants with a 2-day separation between sessions exhibited a similar level of recall compared to those with a 7-day separation. Moreover, this interval has been widely used in prior research on web information re-finding [18, 49, 71, 74].

Two study scenarios were selected for the study. In the first scenario, participants were asked to read the latest finance news and discussions to make an investment decision. In the second scenario,




Fig. 4. The user study procedure. The user study composed of 11 parts: (1) consent form and demographics survey, (2) task 1 for scenario 1, (3) task 2 for scenario 1, (4) task 1 for scenario 2, (5) task 2 for scenario 2, (6) time interval, (7) task 1 reconstruction for scenario 1, (8) task 2 reconstruction for scenario 1, (9) task 1 reconstruction for scenario 2, (10) task 2 reconstruction for scenario 2, (11) semi-structured interview. In the second session, tasks enclosed by dashed lines indicate that participants used IRCHIVER to perform the revisitation task.

they planned a trip to Havana, Cuba, from their current location. The chosen scenarios—online news consumption [3, 72, 96] and trip planning [49, 70, 75]—are commonly used in prior research on information re-finding and online sensemaking. The two scenarios differ in nature: the first scenario required participants to consume and analyze unstructured online information to reach a decision, while the second scenario involved comparing structured dimensions of available options to select the best choice.

Each study scenario consisted of two smaller subtasks. In the investment scenario, the first task required participants to decide on one or more sectors to invest in, informed by the latest Wall Street Journal articles. The second task asked participants to choose one or more specific stocks to buy, based on the latest stock market discussions on Reddit. In the trip planning scenario, the first task involved determining the best plane ticket to the destination, while the second task required selecting the best housing option at the destination. The two tasks within each scenario were designed to be of comparable complexity and contributed to the same overarching context outlined by the scenario.

4.2.3 Pilot Study. We conducted pilot studies with 6 participants to address the following objectives: (1) to determine whether the tasks within each scenario have similar levels of difficulty in terms of reading and cognitive processing effort, (2) to assess whether separating the two sessions by 2 to 7 days affects how participants reconstruct their decision-making process, (3) to identify whether changes to the websites visited during the sensemaking tasks in the first session introduce challenges for revisitation in the second session, and (4) to confirm whether participants can complete both sessions in approximately 60 minutes.

Based on observations of participants' re-finding behaviors, their self-evaluations, and feedback, we confirmed the following: (1) the tasks within each scenario had similar levels of difficulty, (2) participants with a 2-day separation between sessions exhibited similar recall levels to those with a 7-day separation, corroborating theoretical justification above, (3) all participants encountered at least one instance where changes to the initially visited websites made

revisitation challenging, and (4) all participants were able to complete both sessions in approximately 60 minutes. All participants from pilot studies were excluded from the user studies.

4.2.4 First Session—Initial Online Sensemaking. The first session began with obtaining consent and asking participants to fill out a demographic survey. Participants were then instructed to install IRCHIVER on their personal computers. Lastly, participants used their browsers to perform each of the four decision-making tasks. While performing each task, participants were asked to think aloud [92], verbalizing their thought process and how each piece of online information provided them with insights. Each task had a 10-minute time constraint, and participants were allowed to finish a task early if they reached a final decision and believed they had explored most relevant information or options.

One participant had a personal computer that was incompatible with IRCHIVER, so this participant used the researcher's computer with a dedicated Chrome profile set up for this participant. For each participant, the order of the two scenarios and the order of the two smaller tasks within each scenario were counterbalanced. There was no direct interaction between participants and IRCHIVER in this session, as IRCHIVER was solely used to record snapshots and build a personal web archive while participants engaged in initial browsing and decision-making.

Participants were asked to share their screens, and both their screen and audio were recorded during the decision-making tasks. These recordings were then transcribed and analyzed. The primary researcher transcribed each participant's decision-making process during the think-aloud sessions. Each key factor mentioned was paired with the corresponding screenshot from the recorded screen interactions at the moment the participant verbalized it. This artifact was treated as ground-truth data for the second session. Each instance of the first session lasted approximately 60 minutes.

4.2.5 Second Session—Revisit Past Online Sensemaking Tasks. Among the 14 participants who completed the first session, 12 completed the second session. The second session began with obtaining consent. Participants were then instructed to re-open IRCHIVER on their personal computers. Subsequently, participants used their browsers

to perform revisit each of the four sensemaking tasks from the first session. Specifically, they were asked to: (1) re-familiarize themselves with the decision they made in the first session (or where it was left off if incomplete), (2) restore their mental models of the initial decision-making process, including recalling relevant factors that were part of the decision-making, and (3) re-find the original web pages where they had learned about each relevant factor.

In this session, participants were not asked to think aloud; instead, they were instructed to restore their original mental models as fully as possible. For each participant, the order of the two scenarios and the order of the two smaller tasks within each scenario were counterbalanced. For each scenario, participants were asked to perform one underlying task using IRCHIVER (experimental condition) and the other using any other method they deemed reasonable, but not IRCHIVER (control condition). The time each participant spent on each reconstruction task was recorded.

After each participant finished one reconstruction task, they were presented with the ground-truth artifact for that task from the first session, including (1) the decision-making process and (2) screenshots of browser windows from when each relevant factor was mentioned. Participants were then asked to self-evaluate how well they could re-find relevant information, how completely they could restore their mental models, and how accurately they could reconstruct their decision-making process, using a 7-point Likert scale.

Following the revisitation tasks, we conducted a semi-structured interview to ask participants general questions about their experience with the system. Audio and screen recordings were made throughout the entire study session, starting from the completion of the consent form to the end of the semi-structured interview. Finally, participants were instructed to uninstall IRCHIVER from their computers.

5 Results

Through a mixed-methods analysis, we examined the values, behavioral shifts, and real-world implications of IRCHIVER, an information-centric PIM system. We now summarize the findings and offer insights into our research questions on evaluating IRCHIVER's effectiveness in helping people revisit past sensemaking tasks.

5.1 Efficient Information Revisitation through Shortcuts over Complex and Dynamic Navigation Paths

First, participants self-evaluated how well they were able to re-find information in each scenario (Row 1 of Tab. 2). In the investment scenario, participants using IRCHIVER rated their ability to re-find evidence significantly higher (median rating: 7.0) compared to those in the control condition (median rating: 5.5), indicating a notable advantage for IRCHIVER ($p = 0.027$). In contrast, during the trip planning scenario, participants using IRCHIVER also reported higher ratings (median rating: 7.0) than those in the control condition (median rating: 6.0). However, this difference was less prominent and only marginally significant ($p = 0.057$). Notably, there was no significant effect of scenario on perceived verifiability for either the control condition (median ratings of 5.5 in investment vs. 6.0 in trip planning, $p = 0.305$) or IRCHIVER (median rating: 7.0 in both,

$p = 0.739$). This consistency suggests that participants approached both scenarios similarly, but IRCHIVER offered more support in re-finding information, especially in the investment scenario.

Second, participants reported their self-perceived completeness in restoring their mental models (Row 2 of Tab. 2). In the investment scenario, participants using IRCHIVER achieved significantly more complete restoration of the relevant factors integral to their original decision-making process compared to those in the control condition (median rating: 6.0 vs. 4.5, respectively; Wilcoxon signed-rank test: $p = 0.027$). For the trip planning scenario, participants using IRCHIVER also perceived a more complete restoration of the original relevant factors (median rating: 7.0) compared to those in the control condition (median rating: 5.5). However, this difference was less evident and only marginally significant ($p = 0.075$). Notably, there was no significant effect of scenario on perceived completeness for either the control condition (median rating: 4.5 in both scenarios, $p = 0.820$) or IRCHIVER (median rating: 6.0 in both scenarios, $p = 0.942$). This consistency suggests that the difference in perceived completeness was not due to participants approaching the two scenarios differently, but rather that IRCHIVER provided more effective support in the investment scenario than in the trip planning scenario.

From our interviews, many participants (6 out of 12) explained that IRCHIVER offered a shortcut directly to the page and the on-page position where the information resided. The experience of using a full-text search to jump to exactly where the information located was described by P6 as *"one of those magical moments"* and *"this is exactly what I was looking for, so I'm very, very impressed with this one."* Similarly, P4 believed that *"if I just remember the conclusion, I could kind of skip [the intermediate steps]"* and *"fast track"* to the result. This shortcut avoided the cumbersome and time-consuming attempt to replay the original navigation path, including clicking a series of links and applying various filters. For instance, reflecting on the investment scenario, P8 concluded that *"If somehow you need to go to 'About Us' then 'Korea' then 'Market Data', then this 'straight link' may help."* Likewise, P3 believed IRCHIVER made re-finding the flight information *"much faster and easier"* because it avoided the work to *"filter[ing] all the information I want and all the characteristics I want my flight to have."* Moreover, the shortcut became essential for participants due to the dynamic nature of web content. The challenge posed by this dynamic nature is that the navigation path to relevant information can change completely between the initial discovery and subsequent re-finding attempts. According to P3, who was unable to re-find the page that gave insight about a particular financial sector with the control condition, *"there are specific details that stay on my mind, so I'm trying to remember these details, and then I'm trying to remember where I saw these details, and then I try to go to where you saw those details, which is where the problem happens with dynamic information, for example, there's this table on this part of the page,"* and looking at the Wall Street Journal where new content replaced old ones, P3 *"don't remember where that table was, how to get to that."*

This corroborates with the quantitative findings above where, compared to the trip planning scenario, IRCHIVER showcases a more significant effect in the investment scenario, helping participants better re-find where they first saw the information and remember

Tab. 2. Statistics of scores in the post-tasks survey.

Dimension	Question	Investment		Trip Planning	
		Control	IRCHIVER	Control	IRCHIVER
Re-find Information	To what extent do you feel the tool(s) you used helped you re-find the original pages that contain relevant factors in your original decision-making process?	5.5*	7.0*	5.5	6.0
Restore Mental Model	To what extent do you feel the tool(s) you used helped you re-find relevant factors (e.g., specific criteria, pros & cons, trade-offs, etc.) in your original decision-making process?	4.5*	6.0*	4.5	6.0
Reconstruct Decision-Making Process	To what extent do you feel you were able to accurately reconstruct the original decision-making process?	4.5*	6.0*	5.0*	6.5*

relevant factors associated with it. The investment scenario required participants to make decisions based on information reachable after a longer navigational path. Specifically, participants generally learned relevant factors from nested comments of Reddit posts and charts on Wall Street Journal that took multiple clicks to navigate to. On the other hand, in the trip planning task, participants generally found candidate options and relevant factors on the structured list by vendor sites such as Expedia with one query and one click.

5.2 Screenshots as Contextual Anchors for Accurate Reconstruction of Past Sensemaking

Across both scenarios, participants using IRCHIVER achieved significantly more accurate reconstruction of their metacognitive process compared to the control condition, as shown in Row 3 of Tab. 2. In the investment scenario, participants rated their reconstruction accuracy higher with IRCHIVER (median rating: 6.0) compared to the control condition (median rating: 4.5), with a statistically significant difference ($p = 0.017$). Similarly, in the trip planning scenario, participants reported greater accuracy with IRCHIVER (median rating: 6.5) compared to the control condition (median rating: 5.0), which was also statistically significant ($p = 0.024$).

Participants generally (7 out of 12) relied on visual cues from the screenshots to reconstruct their decision-making process. Compared to the built-in browser history, which listed the visited pages and contained “*just factual information*” (P2), the screenshots from IRCHIVER also helped people recall their on-page actions and “*what I did at the moment*” (P5). Moreover, the screenshots helped participants “*trace my thoughts and how I made decisions*” (P1) and invoked participants’ situational context: “*when I saw the screenshot of when I looked up ‘is Miami safe’, I remembered I did this search because I wanted to know if the overnight flight was worth it or not*” (P1). P9 reported seeing a screenshot that showed the mouse hover over the navigation bar on Wall Street Journal brought back the “sense of confusion” felt during the first session where P9 “did not remember which sector to investigate first,” and this was something “*I would not feel by looking at the browser history.*”

Many participants (6 out of 12) liked how IRCHIVER could sort the search results by time because it creates a timeline view for them to “*reconstruct the memory*” (P8) where they “*just need to go through it, and then the whole process will be imaginable again in my head*” (P11). P7 believed the chronological presentation of search results in IRCHIVER was beneficial because “*so that you know from which page you then went to which page*” and “*why you are at that*

page”. Without it, P7 thought “*you lose the context, and you just see a page but you don’t know why you are at the page.*” At the same time, P6 did not like IRCHIVER’s timeline view because it showed intermediary steps along the navigation path such as “*a search engine result page*”, which P6 found to be “*not super interesting.*”

Moreover, participants noted that their decision-making process was often tied to the exact on-page values present at the time of the original decision. Although many participants used the same approach to reconstruct their initial decision, they often failed to reach the same conclusion because key values had changed. For example, when P12 initially booked flights, the itinerary from [the city P12 lives in] to Havana was \$200 more expensive than booking two separate flights from [the city P12 lives in] to Miami and then from Miami to Havana. However, in the second session, a significant price drop for the direct itinerary prevented P12 from reconstructing the original decision. P12 reflected, “*It’s like less than a hundred-dollar difference now, and maybe it’s not worth the hassle of purchasing two flights.*” By preserving information exactly as it was at the time of the decision, IRCHIVER enabled users to access the original values they encountered, aiding them in accurately reconstructing and reflecting on their decision-making process.

5.3 Boosting Confidence by Preserving Information as Originally Seen

After participants finished revisitation tasks for both scenarios in the second session, we asked them to self-evaluate the overall confidence they perceived between the control condition and IRCHIVER. Out of the 12 subjects, 10 agreed that IRCHIVER gave them more confidence, and the other 2 said were indifferent between the two. Participants who felt more confident when using IRCHIVER attributed it to IRCHIVER’s ability to preserve snapshots of the page exactly when they initially saw it. For example, when inspecting resulting screenshots from the query “*healthcare*”, P9 noticed the term “*undervalue*” was used to describe the healthcare sector and was reminded “*that’s why I decided to invest in healthcare sector.*” P9 further explained that “*knowing that I’m seeing exactly what I saw when I made the initial decision, it makes me believe that must be the reason why I made my initial decision.*” According to P8, “*It gives ground truth. With IRCHIVER, it feels like I’m searching among ground truth. With others, it feels like I’m discovering the ground truth [again].*” On the other hand, when revisiting past sensemaking tasks with the control condition, participants were certain that something must have been missed. When describing the mental model reconstruction experience without IRCHIVER, P4 felt “*very*

unconfident [about] the conclusion I was making. I was so sure I forgot something, and I did." In contrast, when using IRCHIVER, P4 felt *"I'm sure I'm not forgetting anything."*

Participants also mentioned their confidence in IRCHIVER was subject to change, depending on whether they could re-find a snapshot for what they have in mind. For example, after using IRCHIVER to retrieve the exact screenshot they had in mind, P6 stated that it *"really boosted my confidence in the search results."* In contrast, P11 experienced a decline in confidence upon realizing that IRCHIVER was *"missing some screenshots for some of the important information."*

5.4 Search-Initiated Mental Model Restoration

When restoring mental models with IRCHIVER, all participants used a search query in IRCHIVER to start the mental model restoration process. Most participants (9 out of 12) indicated they used what they remembered about the task as their first issued query. The first issued query could be something concrete, such as a decided option such as "Nvidia" for P1, because *"that's what I remember, the results."* Alternatively, it could be something more abstract such as "Havana neighborhood" for P6, as they were *"keywords for what I was trying to achieve."* However, P8 expressed concern that *"sometimes it's tricky to remember the keyword,"* and having *"query suggestion"* would be very helpful.

Many participants (5 out of 12) mentioned they strategically entered their queries in a breadth-first manner to get a broad perspective of the original decision-making process. P5 tried a combination of similar queries at the beginning, including "flight to Cuba", "fly to Havana", and "fly to Havana, Cuba" because P5 wanted to *"expand the range as much as possible to remember what exactly I did to come up with the decision, and maybe it gives me a complete picture of my thought process."* According to P4, the reconstruction process had a "tree structure" where the remembered conceptual characteristics about the task, *"I was looking at hotels, [and] I was looking at hostels,"* were the *"roots I can remember and access,"* and that *"from there I can access the branches."*

Visual cues from IRCHIVER's screenshots reminded multiple participants of missing information, prompting them to initiate subsequent search queries. For example, after issuing the first query "Havana Hotel", P9 learned that *"from the hotel search results, there were like a couple of Airbnb pages [screenshots],"* and it led to a subsequent query "Airbnb".

5.5 Enhanced Information Capture without Added Overhead

IRCHIVER would capture multiple snapshots of the same page when participants inspected content spread across different parts of the page, whereas the control condition would capture one single snapshot at most.

To evaluate whether IRCHIVER introduced significant overhead for participants when revisiting past sensemaking tasks, we recorded the time participants spent reconstructing each scenario. In the investment scenario, participants using IRCHIVER took an average of 323 seconds ($\sigma = 183$) to complete the reconstruction compared to 347 seconds ($\sigma = 227$) for those in the control condition. For the trip planning scenario, the average reconstruction time with IRCHIVER was 165 seconds ($\sigma = 101$) compared to 203

seconds ($\sigma = 117$) for the control condition. To further examine these differences, we conducted a repeated-measures ANOVA to test the effect of *Condition* (control vs. IRCHIVER) on reconstruction times for each scenario. Since our data did not follow a normal distribution based on the Shapiro-Wilk test results ($p = 0.027$ for the investment scenario and $p = 0.029$ for the trip planning scenario), we applied a log-transformation to all data points and confirmed normality ($p = 0.974$ for the investment scenario and $p = 0.469$ for the trip planning scenario). The ANOVA results indicated no significant difference in reconstruction times between the control condition and IRCHIVER for both the investment scenario ($F[1, 11] = 0.269, p = 0.614$) and the trip planning scenario ($F[1, 11] = 1.176, p = 0.301$). These findings suggested that IRCHIVER did not introduce significant overhead in revisiting past sensemaking tasks, even while presenting additional information to participants.

5.6 Meeting Diverse Preferences in Past Sensemaking Revisitation

To learn how people would adopt IRCHIVER to their existing workflows, we first tried to understand their current workflows in both information re-finding and mental model restoration, two integral parts in revisiting past online sensemaking tasks. We present the results in Fig. 5.

Participants often used more than one approach to re-find online information where they would first try the top choice and *"fall back"* (P4) to the next approach if the top choice does not work. About half of them preferred re-querying using search engine and about half of the people preferred using browser history. People ordered their priorities based on their experience with each approach. P2 pointed out that search engines was the easiest because it *"show[s] different colors on the things that you clicked"*, and *"those are the indications of how I would go back to the information that I already visited."* When they *"don't know what kind of keyword I input into the Google search"* (P9), they would fall back to browser history. P6 preferred *"send[ing] messages to myself,"* and then later *"try to structure things that I find are still interesting"* while P3 would leave tabs open because it is *"faster than Googling it [again]."*

The participants usually only had one strategy to restore mental models of online sensemaking tasks. The two most frequently used approaches were (1) first re-finding original pages and reconstructing mental models backwards and (2) using only memory. Some participants chose to re-do the original decision-making process in searching for the original mental model. Participants also identified note-taking as a strategy during the initial browsing, especially it is for *serious tasks* (P8).

Overall, when comparing IRCHIVER with their top information re-finding strategies, participants perceived IRCHIVER as a versatile system that was better than all but one existing identified strategy. First, participants believed IRCHIVER would be better than re-querying using search engine when working with dynamic web content (3 out of 12) or when taking a long navigational path towards relevant information (2 out of 12). Second, participants believed IRCHIVER would be better than browser history because the visual cues from the screenshots help them recognize target pages




Fig. 5. Participants usually have multiple strategies and a fallback plan to re-find information where “Search Engine Re-query” and “Browser History” were the most common strategies. Participants usually have one strategy to restore mental models where “Re-finding original pages and reconstructing backwards” and “Memory” were the most common strategies.

(6 out of 12) and provide more context (7 out of 12). Third, participants who used the “messaging to self” strategy also believed a passive system such as IRCHIVER could replace the active curating approach because “*I know that IRCHIVER is gonna have my back*” (P6) so that “*I don’t have to each time take a screenshot myself*” (P2). On the other hand, P3, who identified keeping tabs open as the most used strategy, believed “*it’s still quicker than going to IRCHIVER and opening up the page again*,” especially for frequently-accessed ones. At the same time, P3 agreed that, knowing IRCHIVER was running in the background, it could prompt P3 to close some of the tabs.

Additionally, we learned real-life scenarios in which participants could benefit from IRCHIVER. First, participants (3 out of 12) believed IRCHIVER could be helpful when they needed to re-find online text-based documents because what they remembered was the information on the pages and how the pages look like instead of the titles. Second, participants (3 out of 12), two of whom would actively leave documentations to aid re-finding, admitted that they would not document things while multitasking or under time-pressure. In this case, they would fully rely on IRCHIVER’s automatic curation. Third, participants would use IRCHIVER to track changed information such as price and use the time-series value to inform decisions. For example, IRCHIVER allowed P4 to “*do that analysis on my own*” and P11 to “*see is this price the lowest or like the highest? I can also like, try to find a pattern in those changes, and (...) to predict if the price will change again.*”

For participants to personally adopt IRCHIVER, privacy emerged as the most common concern (8 out of 12). Another shared concern among participants (4 out of 12) was the local disk space required to store the data. P4 specifically emphasized the need for transparency in data usage and an easy method to clear stored data, stating, “*In my browser, I can see exactly how much data it consumes, and I can clear it really easily, so having a feature like that would be nice.*”

6 Discussion

Users agreed that IRCHIVER was helpful and they would incorporate it to their existing workflows. We further elaborate on how a searchable visual personal web archive can support reconstructing past

online sensemaking tasks, including both information re-finding and mental model restoration.

6.1 Information as a First-Class Object Effectively Preserves Content and Mental Models on a Dynamic Web

By treating information as a first-class object, IRCHIVER enabled people to re-find on-page information and restore relevant mental models despite the dynamic navigational paths. Our findings in Sec. 5.1 revealed that the dynamic nature of web content presented a major challenge for users attempting to re-find information. When re-finding information, participants tried to replay the steps they had originally taken to navigate back to specific information they had in mind, using various points, such as a table on the Wall Street Journal for P3, along the navigation path as cues. This strategy aligns with prior studies that found the original small, localized steps were important in later re-finding [100] and that memorable waypoints along the navigation path were helpful during re-finding attempts [19]. However, when previously established waypoints—such as tables or links—were no longer present due to content updates, this approach failed. By creating snapshots and indexing on-page text exactly when people initially saw it, IRCHIVER enabled them to re-find information so they could “*look back on information that I wouldn’t be able to find if the page has changed*” (P4) and restore mental models associated with the information to “*remember the context of what’s going on*” (P12).

By preserving the information exactly as it was consumed, IRCHIVER enabled people to access the original values they initially saw and use them to reconstruct the original decision-making process and eventually reflect on it. In Sec. 5.2, we also found that the dynamic nature of web content complicated users’ attempts to accurately restore and reflect on past decision-making processes. People’s decisions were often influenced by context-specific information available at the time. When critical values or factors changed, the rationale for those original decisions could be lost.

This was confirmed in participants' struggles to reconstruct their decisions among participants in the control condition.

6.2 People Are Ready to Integrate IRCHIVER into Existing Workflows

Overall, we found participants' information re-finding and mental model restoration strategies corroborated with conclusions from prior research more than 15 years ago. We first surveyed participants' strategies in both information re-finding and mental model restorations in Sec. 5.6. The overall status quo for information re-finding strategies aligned with existing literature on information re-access. Jones et al. [58] showed that the most common information re-finding strategy was to "do nothing," where people did not take explicit actions such as bookmarking or taking notes during initial browsing to make re-finding easier, and our findings confirmed that this was still true about 25 years later. Aula et al. found that search engines were the more frequently used approach in re-finding online information, but users often failed to remember the original queries [7]. This was consistent with our findings, and our participants reported that they would fall back to other approaches when they could not recall the search queries. Morris et al. [70] surveyed people's strategies for resuming a suspended investigation in which they reported that 36% of respondents used memory to resume suspended tasks, and 55% of respondents used "Active Storage/Active Retrieval" methods such as written or typed notes and browser bookmarks. While we also observed memory as a popular mental model reconstruction strategy, only a small proportion of participants identified "Active Storage/Active Retrieval" as a strategy they used. This discrepancy likely originated from the subtle difference in our focus; instead of focusing on task resumptions, we asked participants for strategies used to restore mental models in general where they did not know if they would reuse their mental models in the future and therefore would not take notes for their future selves.

Participants, satisfied with IRCHIVER's support for revisiting past sensemaking tasks, believed it offered advantages over existing strategies such as re-querying in search engines, using browser history, and active curation. They also suggested practical use cases, including reducing cluttered browsing and providing universal version control. Participants preferred IRCHIVER over re-querying in search engines when working with dynamic online information or when the navigation path was long and complicated so they could avoid the inefficiencies and frustrations caused by the dynamic nature [71, 89, 98, 105]. Participants liked IRCHIVER more than browser history because the screenshots could help them recognize target pages and evoke more context. This confirmed that the benefits of visual cues in both page recognition [35, 47, 59, 83, 101] and context restoration [49, 86, 87] are transferable to personal web archive systems. Participants often relied on page content and visual layout rather than page titles when trying to re-find text-based documents, which were not supported by browser history tools based solely on titles that often failed to capture the actual content users wanted to revisit [29]. Participants who used active curating strategies such as messaging to themselves were also ready to adopt IRCHIVER over their current strategies because IRCHIVER's continuous curating approach freed them from worrying about

documenting themselves, especially under time pressure or while multitasking. It demonstrated that a continuous passive foraging system could free people from manual active information foraging, which was the most time-consuming phase during a sensemaking process [14, 20, 68, 77]. Although P3 believed revisiting kept-open tabs could still be quicker than re-finding them using IRCHIVER, knowing that IRCHIVER was running in the background prompted P3 to close some of the tabs, showing IRCHIVER's potential to avoid information overload from cluttered browsing [22, 67]. Lastly, participants reported using IRCHIVER to track changed information such as prices and to use the time-series value to inform decisions. Kellar et al. identified monitoring changed information as one motive for information gathering [60], and IRCHIVER showed the potential to become a universal diff tracker for users.

Consistent with previous systems that programmatically curated screenshots [49, 51], privacy was the most common concern among participants regarding the personal adoption of a web archive like IRCHIVER. To address these concerns, IRCHIVER allows users to pause and resume recording and easily directs them to the local directory where screenshots are saved, enabling further inspection and deletion. Additionally, several participants raised concerns about the amount of local disk space required to store all the data. The design of IRCHIVER (Sec. 3) addressed both concerns by using fully local data storage and compressing screenshots into WEBP format, with a more manageable size, automatically.

6.3 Opportunities and Challenges for Personal Web Archives

Based on the findings from the study, we identified the following opportunities and challenges for PIM systems that aid online information re-finding and mental model reconstruction.

First, there was a dichotomy between preserving everything people saw and information overloading. Traditional search engines and browser history did not capture many activities that were meaningful, including (1) switching tabs back and forth (common when cross-validating), (2) reading freshly fetched content by scrolling, and (3) investigating certain parts of the page with actions such as zooming in. While all the above activities were captured by IRCHIVER, it could capture multiple snapshots of the same page, which made people feel confused and eventually led to information overloading [40, 67]. On one hand, it remained important to try to capture everything people saw because they would be frustrated and lose faith in the tool when they could not re-find something they knew they had seen. On the other hand, it was imperative to introduce more advanced aggregation mechanisms to prevent information overloading for long-term usage. Inspired by prior research that clustered browsing experiences by activities and tasks [9, 10, 23, 91], clustering on-page information by tasks appeared to be a promising solution. However, on-page activities and information consumption were more ad-hoc, and parallel browsing [52] added more complexity to the organization and aggregation.

Second, most of the time there was a synergy between re-finding information and reconstructing mental models where one reinforced the other. For example, after issuing a search query in IRCHIVER, participants frequently restored their mental models

while examining results, which often prompted new information re-finding needs and led them to issue additional queries. However, it was difficult to design an interface that promoted both at the same time because they often had different requirements and priorities. When participants tried to re-find online information, they cared more about whether the top results contained the snapshot they had in mind rather than intermediary steps that led to the final page. On the other hand, when participants wanted to reconstruct their mental models, they preferred seeing all search results in chronological order for context restoration, including intermediary snapshots that were part of the navigational path. Given the benefits of search trails presented by prior research [53], one promising design that consolidated two competing revisitation strategies was to prioritize returning the best-matching snapshots from the terminal pages while preserving surrounding context by augmenting the terminal pages with the previous browser trail that led to them and the subsequent browser trail following the inspection of a terminal page.

Third, participants demanded a more personalized ranking algorithm as they did not think the top results presented by IRCHIVER were better than later ones. In addition to document metadata, such as title and published date, Qvarfordt et al. introduced process metadata, such as whether a page had been bookmarked and the number of visits [80]. On top of this, because all the snapshots the ranking algorithm sorted were of pages people had already interacted with and seen at least once, we could further incorporate people's implicit signals during browsing into the ranking algorithm while still promoting the passive nature of IRCHIVER. Specifically, we could use a combination of various implicit signals, including content copying [13, 43, 45], text highlighting [93], clicking [48], cursor movements [25, 41, 48, 53, 85], eye gazing [42, 42, 73], and dwell time [27, 31] to evaluate the relevance of each snapshot [65].

Lastly, while many systems support active curation, few focus on passive curation. When users know what they need to preserve, they can collect resources either through manual foraging tools [66, 88, 94] or automated scraping scripts [24, 26, 33, 57, 78]. However, in unfamiliar domains or open-ended searches, users often cannot anticipate future needs. Under these circumstances, passive curation systems add value by proactively preserving potentially relevant materials without explicit user input. This highlights the need to develop and study passive solutions that safeguard important information over time.

6.4 Limitations and Future Work

The primary aim of this work was to explore the effectiveness and usage of a passive personal web archive in revisiting past online sensemaking tasks in the dynamic web. Designing an optimal search interface was beyond the scope of this study. While many participants appreciated IRCHIVER's lightweight search interface, our findings revealed key design considerations for creating personal web archives that facilitate both information re-finding and mental model restoration. We leave the exploration of these design implications to future research.

Furthermore, with participants expressing interest in incorporating IRCHIVER into their existing workflows, there is an opportunity to expand its deployment. P4 anticipated that additional use cases

would emerge with extended usage, stating, "[I will] find it more useful in the moment, being like 'Oh, I wish I had this.'" To gain deeper insights into its long-term value and usage patterns, we plan to conduct a naturalistic study to observe IRCHIVER's adoption and impact over an extended period. Such a large-scale deployment would also enable us to evaluate potential costs of IRCHIVER, including information overload and memory usage, which are influenced by long-term, real-world usage.

7 Conclusion

This paper explored the value of treating information as a first-class object and preserving dynamic representations of web pages to support users in revisiting past online sensemaking tasks among dynamic web content. We introduced IRCHIVER, a cross-platform system that automatically captures full-resolution screenshots of everything users view in their web browsers, compresses and stores them efficiently, and extracts and indexes on-screen text for easy searching. Through our user study, we found that IRCHIVER enhanced participants' ability to revisit past sensemaking tasks. Specifically, participants were able to re-find information more effectively, restore mental models more completely, and reconstruct decision-making processes with greater accuracy. Participants also reported a boost in confidence during revisitation, attributed to IRCHIVER's ability to preserve exactly what they saw during their original tasks. Additionally, we presented design implications for personal web archives to promote both information re-finding and mental model restoration.

Acknowledgments

We thank Ji Won Chung for feedback on the design and early drafts; Talie Massachi for guidance on participant recruitment; Tongyu Zhou for offering her PC for us to debug; Zainab Iftikhar for IRB suggestions; Brandon Woodard, Joshua Yang, Yizheng Xie, and Eric Chen for feedback on the study design; Ava Hu and Nikos Pagonas for trying an early version of the prototype; and Niki Kittur and Sherry Tongshuang Wu for helpful discussions. We appreciate feedback from the Brown HCI Group. Finally, we are grateful to all our patient participants in both the pilot and user studies.

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, David Karger, and Lynn Andrea Stein. 1999. Haystack: per-user information environments. In *Proceedings of the Eighth International Conference on Information and Knowledge Management* (Kansas City, Missouri, USA) (CIKM '99). Association for Computing Machinery, New York, NY, USA, 413–422. <https://doi.org/10.1145/319950.323231>
- [3] Eytan Adar, Jaime Teevan, and Susan T Dumais. 2008. Large scale analysis of web revisitation patterns. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 1197–1206.
- [4] Scott G. Ainsworth, Ahmed AlSum, Hany SalahEldeen, Michele C. Weigle, and Michael L. Nelson. 2013. How Much of the Web Is Archived? arXiv:1212.6177 [cs.DL] <https://arxiv.org/abs/1212.6177>
- [5] ArchiveBox Team. 2024. ArchiveBox. <https://archivebox.io/>. Accessed: 2024-10-18.
- [6] archive.ph Team. 2024. archive.ph. <https://archive.ph/>. Accessed: 2024-10-18.
- [7] Anne Aula, Natalie Jhaveri, and Mika Käki. 2005. Information search and re-access strategies of experienced web users. In *Proceedings of the 14th international conference on World Wide Web*. 583–592.

- [8] Ziv Bar-Yossef, Andrei Z. Broder, Ravi Kumar, and Andrew Tomkins. 2004. Sic transit gloria telae: towards an understanding of the web's decay. In *Proceedings of the 13th International Conference on World Wide Web* (New York, NY, USA) (WWW '04). Association for Computing Machinery, New York, NY, USA, 328–337. <https://doi.org/10.1145/988672.988716>
- [9] Jakob Bardram, Jonathan Bunde-Pedersen, and Mads Soegaard. 2006. Support for activity-based computing in a personal computing operating system. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*. 211–220.
- [10] Jakob E Bardram, Steven Jeuris, Paolo Tell, Steven Houben, and Stephen Volda. 2019. Activity-centric computing systems. *Commun. ACM* 62, 8 (2019), 72–81.
- [11] Anat Ben-David and Hugo Huurdeman. 2014. Web archive search as research: Methodological and theoretical implications. *Alexandria* 25, 1-2 (2014), 93–111.
- [12] Krishna Bharat. 2000. SearchPad: explicit capture of search context to support Web search. In *Proceedings of the 9th International World Wide Web Conference on Computer Networks: The International Journal of Computer and Telecommunications Networking* (Amsterdam, The Netherlands). North-Holland Publishing Co., NLD, 493–501.
- [13] Joel Brandt, Mira Dontcheva, Marcos Weskamp, and Scott R Klemmer. 2010. Example-centric programming: integrating web search into the development environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 513–522.
- [14] Joel Brandt, Philip J Guo, Joel Lewenstein, Mira Dontcheva, and Scott R Klemmer. 2009. Two studies of opportunistic programming: interleaving web foraging, learning, and writing code. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1589–1598.
- [15] Vannevar Bush. 1945. As We May Think. *The Atlantic Monthly* 176, 1 (July 1945), 101–108.
- [16] Gaston R Cangiano and James D Hollan. 2009. Capturing and restoring the context of everyday work: A case study at a law office. In *Human Centered Design: First International Conference, HCD 2009, Held as Part of HCI International 2009, San Diego, CA, USA, July 19-24, 2009 Proceedings 1*. Springer, 945–954.
- [17] Robert Capra. 2010. HCI browser: A tool for studying web search behavior. *Proceedings of the American Society for Information Science and Technology* 47, 1 (2010), 1–2.
- [18] Robert G. Capra and Manuel A. Pérez-Quinones. 2003. Re-Finding Found Things: An Exploratory Study of How Users Re-Find Information. *ArXiv cs.HC/0310011* (2003). <https://api.semanticscholar.org/CorpusID:14057639>
- [19] Robert G. Capra and Manuel A. Pérez-Quinones. 2005. Using Web search engines to find and refine information. *Computer* 38 (2005), 36–42. <https://api.semanticscholar.org/CorpusID:88050>
- [20] Nicholas J Cepeda, Harold Pashler, Edward Vul, John T Wixted, and Doug Rohrer. 2006. Distributed practice in verbal recall tasks: A review and quantitative synthesis. *Psychological bulletin* 132, 3 (2006), 354.
- [21] Soumen Chakrabarti, Sandeep Srivastava, Mallela Subramanyam, and Mitul Tiwari. 2000. Using Memex to archive and mine community Web browsing experience. *Computer Networks* 33, 1-6 (2000), 669–684.
- [22] Joseph Chee Chang, Nathan Hahn, Yongsung Kim, Juliana Coupland, Bradley Breneisen, Hannah S Kim, John Hwang, and Aniket Kittur. 2021. When the Tab Comes Due: Challenges in the Cost Structure of Browser Tab Usage. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 148, 15 pages. <https://doi.org/10.1145/3411764.3445585>
- [23] Joseph Chee Chang, Yongsung Kim, Victor Miller, Michael Xieyang Liu, Brad A Myers, and Aniket Kittur. 2021. Tabs. do: Task-centric browser tab management. In *The 34th Annual ACM Symposium on User Interface Software and Technology*. 663–676.
- [24] Sarah E Chasins, Maria Mueller, and Rastislav Bodik. 2018. Rousillon: Scraping distributed hierarchical web data. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. 963–975.
- [25] Mon Chu Chen, John R Anderson, and Myeong Ho Sohn. 2001. What can a mouse cursor tell us more? Correlation of eye/mouse movements on web browsing. In *CHI'01 extended abstracts on Human factors in computing systems*. 281–282.
- [26] Weihao Chen, Xiaoyu Liu, Jiacheng Zhang, Ian Iong Lam, Zhicheng Huang, Rui Dong, Xinyu Wang, and Tianyi Zhang. 2023. MIWA: Mixed-Initiative Web Automation for Better User Control and Confidence. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–15.
- [27] Mark Claypool, Phong Le, Makoto Wased, and David Brown. 2001. Implicit interest indicators. In *Proceedings of the 6th international conference on Intelligent user interfaces*. 33–40.
- [28] Paul Clough and Bettina Berendt. 2009. Report on the TrebleCLEF query log analysis workshop 2009. In *ACM SIGIR Forum*, Vol. 43. ACM, New York, NY, 71–77. Issue 2.
- [29] Andy Cockburn and Bruce McKenzie. 2001. What do web users do? An empirical analysis of web use. *Int. J. Hum.-Comput. Stud.* 54, 6 (jun 2001), 903–922. <https://doi.org/10.1006/ijhc.2001.0459>
- [30] Edward Cutrell, Daniel Robbins, Susan Dumais, and Raman Sarin. 2006. Fast, flexible filtering with phlat. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Montréal, Québec, Canada) (CHI '06). Association for Computing Machinery, New York, NY, USA, 261–270. <https://doi.org/10.1145/1124772.1124812>
- [31] Sergio Cozzetti B de Souza, Nicolas Anquetil, and Káthia M de Oliveira. 2005. A study of the documentation essential to software maintenance. In *Proceedings of the 23rd annual international conference on Design of communication: documenting & designing for pervasive information*. 68–75.
- [32] Tangjian Deng, Liang Zhao, Ling Feng, and Wenwei Xue. 2011. Information re-finding by context: a brain memory inspired approach. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management* (Glasgow, Scotland, UK) (CIKM '11). Association for Computing Machinery, New York, NY, USA, 1553–1558. <https://doi.org/10.1145/2063576.2063799>
- [33] Rui Dong, Zhicheng Huang, Ian Iong Lam, Yan Chen, and Xinyu Wang. 2022. WebRobot: web robotic process automation using interactive programming-by-demonstration. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. 152–167.
- [34] Patrick Dubroy and Ravin Balakrishnan. 2010. A study of tabbed browsing among mozilla firefox users. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 673–682.
- [35] Susan Dumais, Edward Cutrell, JJ Cadiz, Gavin Jancke, Raman Sarin, and Daniel C. Robbins. 2003. Stuff I've seen: a system for personal information retrieval and re-use. In *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Informaion Retrieval* (Toronto, Canada) (SIGIR '03). Association for Computing Machinery, New York, NY, USA, 72–79. <https://doi.org/10.1145/860435.860451>
- [36] H. Ebbinghaus. 1913. *Memory: A Contribution to Experimental Psychology*. Teachers College, Columbia University, New York. Originally published in 1885.
- [37] David Elweiler, Mark Baillie, and Ian Ruthven. 2008. Exploring memory in email refinding. *ACM Trans. Inf. Syst.* 26, 4, Article 21 (oct 2008), 36 pages. <https://doi.org/10.1145/1402256.1402260>
- [38] Henry A Feild and James Allan. 2013. Using CrowdLogger for in situ information retrieval system evaluation. In *Proceedings of the 2013 workshop on living labs for information retrieval evaluation*. ACM, New York, NY, 15–18.
- [39] Jim Gemmell, Gordon Bell, Roger Lueder, Steven Drucker, and Curtis Wong. 2002. MyLifeBits: fulfilling the Memex vision. In *Proceedings of the Tenth ACM International Conference on Multimedia* (Juan-les-Pins, France) (MULTIMEDIA '02). Association for Computing Machinery, New York, NY, USA, 235–238. <https://doi.org/10.1145/641007.641053>
- [40] Benedikt Graf and Conny H Antoni. 2022. Drowning in the flood of information: A meta-analysis on the relation between information overload, behaviour, experience, and health and moderating factors. *European Journal of Work and Organizational Psychology* 6, 1 (2022), 1–26.
- [41] Qi Guo and Eugene Agichtein. 2010. Towards predicting web searcher gaze position from mouse movements. In *CHI'10 extended abstracts on human factors in computing systems*. 3601–3606.
- [42] Jacek Gwizdzka. 2014. Characterizing relevance with eye-tracking measures. In *Proceedings of the 5th information interaction in context symposium*. 58–67.
- [43] Björn Hartmann, Mark Dhillon, and Matthew K Chan. 2011. HyperSource: bridging the gap between source and code-related web sites. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2207–2210.
- [44] Koichi Hayashi, Takahiko Nomura, Tan Hazama, Makoto Takeoka, Sunao Hashimoto, and Stephan Gumundson. 1998. Temporally threaded workspace: a model for providing activity-based perspectives on document spaces. In *Proceedings of the ninth ACM conference on Hypertext and hypermedia: links, objects, time and space—structure in hypermedia systems: links, objects, time and space—structure in hypermedia systems*. 87–96.
- [45] Andrew Head, Elena L Glassman, Björn Hartmann, and Marti A Hearst. 2018. Interactive extraction of examples from existing code. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [46] Ron R. Hightower, Laura T. Ring, Jonathan I. Helfman, Benjamin B. Bederson, and James Hollan. 1998. Graphical multiscale Web histories: a study of padprints. In *UK Conference on Hypertext*. <https://api.semanticscholar.org/CorpusID:2891425>
- [47] Ron R. Hightower, Laura T. Ring, Jonathan I. Helfman, Benjamin B. Bederson, and James D. Hollan. 1998. Graphical multiscale Web histories: a study of padprints. In *Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia: Links, Objects, Time and Space—Structure in Hypermedia Systems: Links, Objects, Time and Space—Structure in Hypermedia Systems* (Pittsburgh, Pennsylvania, USA) (HYPERTEXT '98). Association for Computing Machinery, New York, NY, USA, 58–65. <https://doi.org/10.1145/276627.276634>
- [48] Yoshinori Hijikata. 2004. Implicit user profiling for on demand relevance feedback. In *Proceedings of the 9th international conference on Intelligent user interfaces*. 198–205.
- [49] Donghan Hu and Sang Won Lee. 2020. ScreenTrack: Using a Visual History of a Computer Screen to Retrieve Documents and Web Pages. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu,

- HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376753>
- [50] Donghan Hu and Sang Won Lee. 2022. Scrapbook: Screenshot-Based Bookmarks for Effective Digital Resource Curation across Applications. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) (UIST '22). Association for Computing Machinery, New York, NY, USA, Article 96, 13 pages. <https://doi.org/10.1145/3526113.3545678>
- [51] Donghan Hu and Sang Won Lee. 2024. Exploring the Effectiveness of Time-lapse Screen Recording for Self-Reflection in Work Context. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–14.
- [52] Jeff Huang and Ryen W White. 2010. Parallel browsing behavior on the web. In *Proceedings of the 21st ACM conference on Hypertext and Hypermedia*. 13–18.
- [53] Jeff Huang, Ryen W White, Georg Buscher, and Kuansan Wang. 2012. Improving searcher models using mouse cursor activity. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*. 195–204.
- [54] Internet Archive. 2024. Internet Archive. <https://archive.org/>. Accessed: 2024-10-18.
- [55] Internet Archive. 2024. Wayback Machine APIs. https://archive.org/help/wayback_api.php. Accessed: 2024-08-15.
- [56] Farnaz Jahanbakhsh, Ahmed Hassan Awadallah, Susan T Dumais, and Xuhai Xu. 2020. Effects of past interactions on user experience with recommended documents. In *Proceedings of the 2020 Conference on Human Information Interaction and Retrieval*. 153–162.
- [57] Dhanya Jayagopal, Justin Lubin, and Sarah E Chasins. 2022. Exploring the learnability of program synthesizers by novice programmers. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–15.
- [58] William Jones, Harry Bruce, and Susan Dumais. 2001. Keeping found things found on the web. In *Proceedings of the Tenth International Conference on Information and Knowledge Management* (Atlanta, Georgia, USA) (CIKM '01). Association for Computing Machinery, New York, NY, USA, 119–126. <https://doi.org/10.1145/502585.502607>
- [59] Shaun Kaasten, Saul Greenberg, and Christopher Edwards. 2002. How People Recognise Previously Seen Web Pages from Titles, URLs and Thumbnails. In *People and Computers XVI - Memorable Yet Invisible*, Kristine Faulkner, Janet Finlay, and Françoise Détéienne (Eds.). Springer London, London, 247–265.
- [60] Melanie Kellar, Carolyn Watters, and Kori M Inkpen. 2007. An exploration of web-based monitoring: implications for design. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 377–386.
- [61] Johannes Kiesel, Arjen P de Vries, Matthias Hagen, Benno Stein, and Martin Potthast. 2018. WASP: web archiving and search personalized. In *Proceedings of the First Biennial Conference on Design of Experimental Search & Information Retrieval Systems*. Ceur Workshop Proceedings, Bertinoro, Italy, 16–21.
- [62] Johannes Kiesel, Florian Kneist, Milad Alshomary, Benno Stein, Matthias Hagen, and Martin Potthast. 2018. Reproducible web corpora: interactive archiving with automatic quality assessment. *Journal of Data and Information Quality* (JDIQ) 10, 4 (2018), 1–25.
- [63] Aniket Kittur, Andrew M Peters, Abdigani Diriyee, Trupti Telang, and Michael R Bove. 2013. Costs and benefits of structured information foraging. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2989–2998.
- [64] 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*. 67–80.
- [65] Michael Xieyang Liu, Aniket Kittur, and Brad A Myers. 2022. Crystalline: Lowering the Cost for Developers to Collect and Organize Information for Decision Making. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [66] Michael Xieyang Liu, Andrew Kuznetsov, Yongsung Kim, Joseph Chee Chang, Aniket Kittur, and Brad A Myers. 2022. Wigglyte: Low-cost information collection and triage. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–16.
- [67] Rongjun Ma, Henrik Lassila, Leysan Nurgalieva, and Janne Lindqvist. 2023. When Browsing Gets Cluttered: Exploring and Modeling Interactions of Browsing Clutter, Browsing Habits, and Coping. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 861, 29 pages. <https://doi.org/10.1145/3544548.3580690>
- [68] Gary Marchionini. 1995. *Information seeking in electronic environments*. Cambridge university press.
- [69] John Markwell and David W. Brooks. 2003. Link rot limits the usefulness of web-based educational materials in biochemistry and molecular biology. *Biochemistry and Molecular Biology Education* 31, 1 (2003), 69–72. <https://doi.org/10.1002/bmb.2003.31.1.69>
- [70] Dan Morris, Meredith Ringel Morris, and Gina Venolia. 2008. SearchBar: a search-centric web history for task resumption and information re-finding. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy) (CHI '08). Association for Computing Machinery, New York, NY, USA, 1207–1216. <https://doi.org/10.1145/1357054.1357242>
- [71] Hartmut Obendorf, Harald Weinreich, Eelco Herder, and Matthias Mayer. 2007. Web page revisitation revisited: implications of a long-term click-stream study of browser usage. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '07). Association for Computing Machinery, New York, NY, USA, 597–606. <https://doi.org/10.1145/1240624.1240719>
- [72] Shuya Ochiai, Makoto P Kato, and Katsumi Tanaka. 2014. Re-call and re-cognition in episode re-retrieval: A user study on news re-finding a fortnight later. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*. 579–588.
- [73] Alexandra Papoutsaki, James Laskey, and Jeff Huang. 2017. Searchgazer: Webcam eye tracking for remote studies of web search. In *Proceedings of the 2017 conference on conference human information interaction and retrieval*. 17–26.
- [74] Haekyu Park, Gonzalo Ramos, Jina Suh, Christopher Meek, Rachel Ng, and Mary Czerwinski. 2023. FoundWright: A System to Help People Re-find Pages from Their Web-history. arXiv:2305.07930 [cs.HC] <https://arxiv.org/abs/2305.07930>
- [75] Sharoda A Paul and Meredith Ringel Morris. 2009. CoSense: enhancing sensemaking for collaborative web search. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1771–1780.
- [76] Zeynep Pehlivan, Myriam Ben-Saad, and Stéphane Gançarski. 2010. Vi-DIFF: Understanding web pages changes. In *Database and Expert Systems Applications: 21st International Conference, DEXA 2010, Bilbao, Spain, August 30-September 3, 2010, Proceedings, Part I 21*. Springer, 1–15.
- [77] Peter Pirolli and Stuart Card. 1999. Information foraging. *Psychological review* 106, 4 (1999), 643.
- [78] Kevin Pu, Rainey Fu, Rui Dong, Xinyu Wang, Yan Chen, and Toví Grossman. 2022. Semantic: Specifying content-based semantic conditions for web automation programs. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–16.
- [79] Jonathan Pulliza and Chirag Shah. 2018. Information Retrieval and Interaction System (IRIS) A Toolkit for Investigating Information Retrieval and Interaction Activities. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval*. ACM, New York, NY, 333–335.
- [80] Pernilla Qvarfordt, Simon Tretter, Gene Golovchinsky, and Tony Dunnigan. 2014. SearchPanel: framing complex search needs. In *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval* (Gold Coast, Queensland, Australia) (SIGIR '14). Association for Computing Machinery, New York, NY, USA, 495–504. <https://doi.org/10.1145/2600428.2609620>
- [81] Jun Rekimoto. 1999. Time-machine computing: a time-centric approach for the information environment. In *Proceedings of the 12th annual ACM symposium on User interface software and technology*. 45–54.
- [82] Sarah Rhodes. 2010. Breaking Down Link Rot: The Chesapeake Project Legal Information Archive's Examination of URL Stability. *Digital Preservation Publications* 102 (09 2010).
- [83] 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>
- [84] Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gattford, et al. 1995. Okapi at TREC-3. *Nist Special Publication Sp* 109 (1995), 109.
- [85] Kerry Rodden, Xin Fu, Anne Aula, and Ian Spiro. 2008. Eye-mouse coordination patterns on web search results pages. *CHI '08 Extended Abstracts on Human Factors in Computing Systems* (2008). <https://api.semanticscholar.org/CorpusID:1759484>
- [86] Adam Rule, Karen Boyd, Jim Hollan, and Aurélien Tabard. 2015. Restoring the context of interrupted work with desktop thumbnails. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 37.
- [87] Adam Rule, Aurélien Tabard, and Jim Hollan. 2017. Using visual histories to reconstruct the mental context of suspended activities. *Human-Computer Interaction* 32, 5-6 (2017), 511–558.
- [88] MC 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*. 172–181.
- [89] Erik Selberg and Oren Etzioni. 2000. On the instability of web search engines. In *Content-Based Multimedia Information Access - Volume 1* (Paris, France) (RIA0 '00). LE CENTRE DE HAUTES ETUDES INTERNATIONALES D'INFORMATIQUE DOCUMENTAIRE, Paris, FRA, 223–236.
- [90] Carmine Sellitto. 2005. The impact of impermanent Web-located citations: A study of 123 scholarly conference publications. *J. Assoc. Inf. Sci. Technol.* 56 (2005), 695–703. <https://api.semanticscholar.org/CorpusID:12391009>
- [91] Chirag Shah, Ryen White, Paul Thomas, Bhaskar Mitra, Shawon Sarkar, and Nicholas Belkin. 2023. Taking search to task. In *Proceedings of the 2023 Conference*

- on *Human Information Interaction and Retrieval*. 1–13.
- [92] Maarten Someren, Yvonne Barnard, and Jacobijn Sandberg. 1994. *The Think Aloud Method - A Practical Guide to Modelling Cognitive Processes*.
 - [93] Jeffrey Stylos and Brad A Myers. 2006. Mica: A web-search tool for finding api components and examples. In *Visual Languages and Human-Centric Computing (VL/HCC'06)*. IEEE, 195–202.
 - [94] Jeffrey Stylos, Brad A Myers, and Andrew Faulring. 2004. Citrine: providing intelligent copy-and-paste. In *Proceedings of the 17th annual ACM symposium on User interface software and technology*. 185–188.
 - [95] Linda Tauscher and Saul Greenberg. 1997. Revisitation patterns in World Wide Web navigation. 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, 399–406. <https://doi.org/10.1145/258549.258816>
 - [96] Jaime Teevan. 2004. How People Re-find Information When the Web Changes. <https://api.semanticscholar.org/CorpusID:12556211>
 - [97] Jaime Teevan. 2007. The re:search engine: simultaneous support for finding and re-finding. In *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology* (Newport, Rhode Island, USA) (*UIST '07*). Association for Computing Machinery, New York, NY, USA, 23–32. <https://doi.org/10.1145/1294211.1294217>
 - [98] Jaime Teevan, Eytan Adar, Rosie Jones, and Michael Potts. 2006. History repeats itself: repeat queries in Yahoo's logs. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*. 703–704.
 - [99] Jaime Teevan, Eytan Adar, Rosie Jones, and Michael A. S. Potts. 2007. Information re-retrieval: repeat queries in Yahoo's logs. In *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Amsterdam, The Netherlands) (*SIGIR '07*). Association for Computing Machinery, New York, NY, USA, 151–158. <https://doi.org/10.1145/1277741.1277770>
 - [100] Jaime Teevan, Christine Alvarado, Mark S. Ackerman, and David R. Karger. 2004. The perfect search engine is not enough: a study of orienteering behavior in directed search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vienna, Austria) (*CHI '04*). Association for Computing Machinery, New York, NY, USA, 415–422. <https://doi.org/10.1145/985692.985745>
 - [101] Jaime Teevan, Edward Cutrell, Danyel Fisher, Steven M. Drucker, Gonzalo Ramos, Paul André, and Chang Hu. 2009. Visual snippets: summarizing web pages for search and revisitation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, MA, USA) (*CHI '09*). Association for Computing Machinery, New York, NY, USA, 2023–2032. <https://doi.org/10.1145/1518701.1519008>
 - [102] Webrecorder Team. 2024. Webrecorder. <https://webrecorder.net/>. Accessed: 2024-10-18.
 - [103] 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 (mar 2008), 31 pages. <https://doi.org/10.1145/1326561.1326566>
 - [104] James Wen. 2003. Post-Valued Recall Web Pages: User Disorientation Hits the Big Time. *IT & Society* 1, 3 (Winter 2003), 184–194. <http://www.ITandSociety.org>
 - [105] Ryan W White, Ian Ruthven, and Joemon M Jose. 2002. Finding relevant documents using top ranking sentences: an evaluation of two alternative schemes. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*. 57–64.
 - [106] Wikipedia contributors. 2024. Category: Articles with permanently dead external links. https://en.wikipedia.org/wiki/Category:Articles_with_permanently_dead_external_links. Accessed: 2024-08-15.