Patterns of Hypertext-Augmented Sensemaking

Siyi Zhu
University of Maryland
United States
zhusy@umd.edu

Robert Haisfield
WebSim
United States
hello@robhaisfield.com

Brendan Langen
Heyday
United States
brendan.langen@gmail.com

Joel Chan
University of Maryland
United States
joelchan@umd.edu

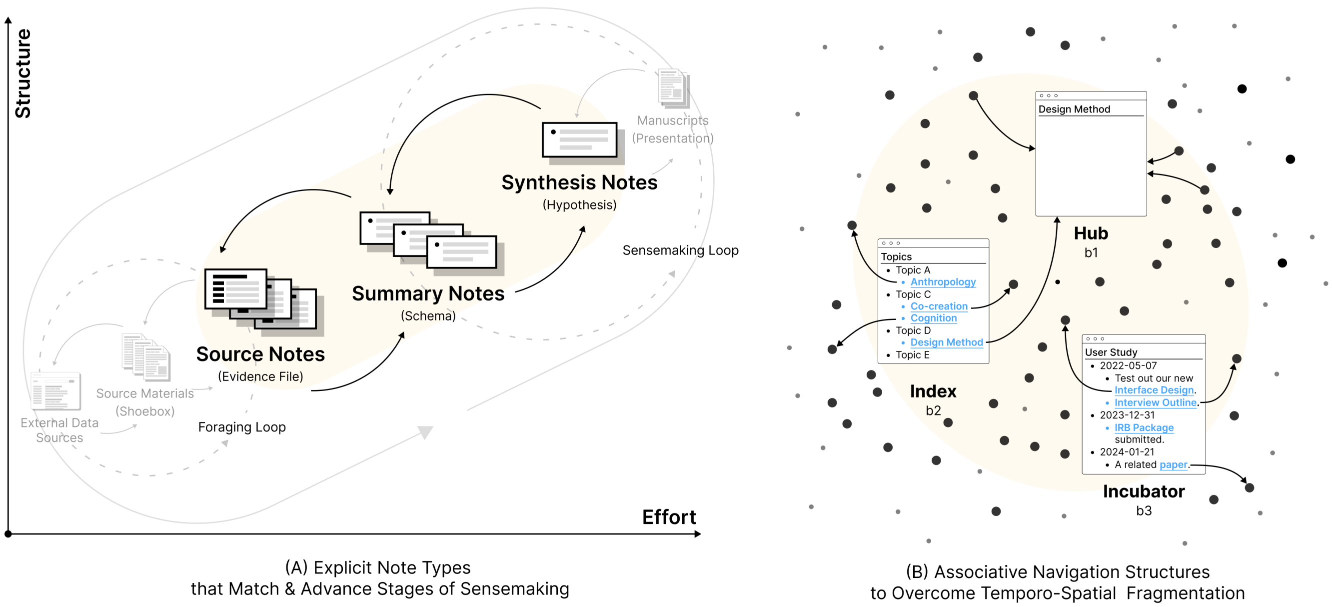

Figure 1: Patterns of hypertext-augmented sensemaking identified in our study: (A) explicit note types that match and advance stages of sensemaking, and (B) associative navigation structures that overcome temporo-spatial information fragmentation.

ABSTRACT

The early days of HCI were marked by bold visions of hypertext as a transformative medium for augmented sensemaking, exemplified in systems like Memex, Xanadu, and NoteCards. Today, however, hypertext is often disconnected from discussions of the future of sensemaking. In this paper, we investigate how the recent resurgence in hypertext "tools for thought" might point to new directions for hypertext-augmented sensemaking. Drawing on detailed analyses of guided tours with 23 scholars, we describe hypertext-augmented use patterns for dealing with the core problem of revisiting and reusing existing/past ideas during scholarly sensemaking. We then discuss how these use patterns validate and extend existing knowledge of hypertext design patterns for sensemaking, and point to new design opportunities for augmented sensemaking.

UIST '24, October 13–16, 2024, Pittsburgh, PA, USA

ACM ISBN 979-8-4007-0628-8/24/10
https://doi.org/10.1145/3654777.3676338

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**; **Hypertext / hypermedia**.

KEYWORDS

sensemaking, hypertext, scholarly synthesis

ACM Reference Format:
Siyi Zhu, Robert Haisfield, Brendan Langen, and Joel Chan. 2024. Patterns of Hypertext-Augmented Sensemaking. In *The 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24), October 13–16, 2024, Pittsburgh, PA, USA.* ACM, New York, NY, USA, 17 pages. https://doi.org/10.1145/3654777.3676338

1 INTRODUCTION

In this paper, we explore the potential of **augmented sensemaking with hypertext**. *Sensemaking* is the process of developing a schema or representation of a problem or situation (often based on data and materials) in order to make a decision or solve a problem [44, 48, 62]. Sensemaking is central to many human activities, from everyday decision-making, to complex problem solving and creative activities such as design, scientific discovery, and policy-making. *Hypertext* is a technological medium that emphasizes a) curating/developing atomic chunks of information, and b) *linking* and *composing* these atomic chunks into larger information structures (see [4] for a history of hypertext as a medium).

The early days of HCI were marked by bold visions of hypertext as a transformative medium for augmented sensemaking. Consider, for instance, Nelson's [39] proposal for a hypertext file structure to enable people to grapple with "the complex, the changing and the indeterminate"; Engelbart's [19] vision for hypertext as a medium for augmenting (collective) human intellect; the gIBIS hypertext tool for grappling with "wicked" policy problems [13]; and Halasz's reflections on authentic usage of the NoteCards system for augmenting knowledge work [25].

Today, however, hypertext is often disconnected from discussions of the future of sensemaking. For example, hypertext is often conflated with the World Wide Web [47]. On the Web, instead of the more granular atomic chunks of information, and structured, bi-directional links and compositional structures such as embedding and transclusion, the primary node is the coarse-grained "document", and links are one way[1] and do not describe the nature of the link. Wikipedia — another prominent manifestation of hypertext — is also quite document-centric over a single view of the hypertext, with limited usage of more advanced hypertext concepts like bi-directional linking and transclusion [4]. When hypertext does come up in relation to sensemaking, it is perhaps more common to emphasize its pitfalls, such as providing more powerful vectors for misinformation on social media.

Does the disconnect between the early visions of hypertext and the discourse on augmented sensemaking today suggest that explorations of the future of sensemaking focus more on novel technologies like AI and extended reality, and leave hypertext as a historical footnote? For example, do the early commercial failures of hypertext systems like *Xanadu* [60], or the shift from repository models to "people-centric" knowledge and expertise sharing systems in CSCW [1], signal some fundamental flaw in the potential of hypertext to augment sensemaking? Are the key open problems for hypermedia described by Halasz et al [25] intractable?

Over the past 5 years, the rise of a new wave of consumer hypertext note-taking tools like Roam Research[2], Obsidian [3], Notion[4], and Tana[5] have brought these questions into sharper focus. Do these tools represent progress on some of the open problems from the first wave of widely adopted hypertext tools [25]? Do use patterns with these new tools suggest interesting new paths to the future of sensemaking that involve hypertext?

We explore these questions with detailed analyses of guided tours from people who use hypertext notebooks for **scholarly sensemaking** (i.e., sensemaking where the core inputs are research papers/ideas). We share specific use patterns that our participants developed to deal with the core problem of *temporo-spatial fragmentation of information* during scholarly sensemaking: 1) define explicit note types that match and progress through stages of their sensemaking process, and 2) flexible and powerful structures for associative navigation, such as hubs, indices, and incubators. These patterns are built from classic hypertext primitives like granular nodes and bi-directional links, but also leverage newer ones like structure search (i.e., specifying queries over a hypertext network based on both content, metadata, and the structure of the hypertext network), and mixed-initiative AI workflows that leverage and contribute to participants' hypertext structures.

Our results validate and extend existing knowledge of hypertext design patterns for sensemaking. For instance, our participants' indices and incubators enable more flexible linkage structures than the strict hierarchical overviews of previous hypertext tools; and their inclusion of structure searches in indices and incubators advance key problems around searching and querying in hypermedia networks. Our findings also reveal future directions for hypertext-augmented sensemaking, such as the need for more powerful end-user programming [33] and mixed-initiative [30] mechanisms to ease the cost of creating and using hypertext structures.

The chief contribution of this paper to the UIST community lies in the detailed descriptions of new design and usage patterns for complex sensemaking work in hypertext, a medium not often studied in HCI and UIST for sensemaking, and reflections on how they might advance our capacity to solve longstanding HCI problems in sensemaking and personal information management.

2 RELATED WORK

Our work is grounded in and extends the conversation in HCI on the use of hypertext for augmented sensemaking. Because our results focus on how users leverage hypertext to address the temporal and spatial fragmentation of information during sensemaking, our work also draws from and extends ideas in the area of personal information management.

2.1 Models and Tools for Sensemaking

We situate our research in HCI models of sensemaking, such as Russell et al's cost structure model [48], and Pirolli et al's Notional Model of Sensemaking [44].

Two core insights from these models inform our work. First, requirements for the nature and extent of tooling support for sensemaking vary by the particulars of the task [23, 62]. For example, depending on the task, structuring representations such as hierarchies, facets, timelines, and networks might be needed [23]. In this work, we are particularly interested in useful representations and tools for complex, open-ended sensemaking, such as scholarly synthesis [46] of prior literature to formulate new research problems or theories [53]. Second, sensemaking work heavily involves *foraging*

[1]Unless you pay for services that reveal backlinks to your website, or are technically savvy enough to implement "indieweb" protocols like Webmentions
[2]https://roamresearch.com/
[3]https://obsidian.md/
[4]https://www.notion.so/
[5]https://tana.inc/

for information [44, 48]. This foraging work involves retrieving external sources of information to kickstart the sensemaking process. But the process of sensemaking itself iteratively shapes and is shaped by foraging: as the sensemaker constructs and develops their overall schema, they frequently loop back over previously retrieved information materials, reflect on how they compare to their schema, and integrate them into their schema. Material that does not "fit" a current schema — called "residue" in Russell's Learning Loop Complex model [48] — can spur reflection and revisions of the schema. Importantly, this foraging process operates not only over external resources (e.g., in a search engine over the Web), but also over a range of previously found material that have gone through varying levels of processing and structuring, such as highlights and annotations on previously read papers.

HCI research has explored a range of designs for supporting sensemaking. For example, researchers have explored affordances for comparing and adding structure to data in spatial canvases [54], structured tables integrated into a Web browser [9], and graph-based visualizations [12]. Other research has explored supporting foraging amongst external sources, integrated browsers extensions for structuring and annotating web sources during sensemaking [34], leveraging information from existing paper collections and notes to suggest web queries [43] or structure search results and recommendations from external sources [10, 31]. Closer to our focus, some past work has also explored how to support re-finding, reuse, and resumption during sensemaking. Some solution patterns include enabling users to add lightweight and flexible annotations and structure to information materials to aid reuse for downstream sensemaking [8, 29, 35], automatically including various levels of context and provenance for information materials in order to facilitate reuse [27], or recommending previously encountered materials, such as prior analyses [40] relevant to current sensemaking tasks.

Our work advances HCI research on sensemaking by connecting the concerns of foraging in sensemaking with problems and solutions in personal information management, as well as with empirical descriptions of "in vivo" use patterns in novel hypertext tools for refinding and reusing information during sensemaking.

2.2 Personal Information Management

Because sensemaking involves foraging over internal sources, problems of refinding and reusing information overlap with problems in personal information management and knowledge management. Much of this literature focuses on how people manage emails and tasks, personal photos and music, and various documents [5, 57]. There is also a robust thread of HCI research on the information management practices of "knowledge workers" [32] — including those doing open-ended complex sensemaking such as designers [15, 17, 51] and scientists [27, 41, 42, 55]

Our work here connects with two core themes from this literature. First, information is typically fragmented across many different tools and devices [6, 24, 55], making it difficult to (re)find needed information [27, 51]. Second, typical organizational structures in digital tools — folders and search — are often insufficient solutions for the fragmentation problem [24]. Many people prefer incremental "orienteering" strategies — navigating towards their information goal in small steps, adjusting/narrowing their query based on contextual information at each step — instead of direct (keyword) search [56]. Folders and categorization systems can also be burdensome to maintain or conflict with personal styles [36]. The rigid hierarchies of folders may also conflict with the open-ended nature of complex sensemaking work [24], or the need for multiple forms of organization, such as topical and chronological views [41]: in such cases, efforts to structure and organize information may actually impede — rather than augment — sensemaking. Yet, a lack of structure altogether impedes refinding and reusing needed information [27, 51]. The dream for truly flexible organization systems for complex sensemaking is not yet reached.

Here, we show how users construct sophisticated structures for refinding and reusing information without complex hierarchies or extensive search, instead leveraging hypertext to support flexible "orienteering" [56] through previously encountered information during sensemaking. Our descriptions also complement technical HCI work on methods for *link-based* approaches to refinding and reusing personal information, such as surfacing and visualizing logical dependencies between files [24].

2.3 Hypertext and Augmented Sensemaking

Our work is rooted in original visions of augmented sensemaking with hypertext, exemplified in proposals and tools like Xanadu [39]; Engelbart's [19] vision for hypertext as a medium for augmenting (collective) human intellect; and spatial and argumentation-based hypermedia tools for grappling with "wicked" policy problems, such as gIBIS[13]. Early waves of hypertext tools — some from academia, such as NoteCards [26] and HyperTies [52], and some from industry, such as Apple's HyperCard — saw authentic, extended usage for sensemaking work and personal information management, resulting in rich empirical descriptions of long-term usage patterns [20, 25, 59]. Reflecting on these long-term deployments, Halasz [25] described key open issues for the next generation of hypermedia systems, including the need for more dynamic (vs. static) data structures, enhanced search and computation over the network, and enhanced collaboration.

More recent work in hypertext — particularly in HCI — has departed from these early visions. An increasingly large proportion of research on hypertext now focuses on the World Wide Web (and related technologies, such as social media) [3, 47]. In partial recognition of this trend, the flagship ACM Conference on Hypertext was renamed in 2012 to the ACM Conference on Hypertext *and Social Media* (emphasis ours) [3].

Separate from the research literature on hypertext, our work responds to the emergence of a new wave of general-purpose "hypertext notebooks" over the past few years. These notebooks implement core hypertext primitives such as granular subdocuments, and bi-directional hyperlinking as a central organizing principle, as opposed to top-down hierarchies of folders and tags. There are also some interesting differences in affordances and values of these communities: for example, several tools, such as Roam Research[6], Logseq[7], and Tana[8], implement hypertext in an "outliner" format, where individual bullets are uniquely addressable subdocuments

[6]https://roamresearch.com/
[7]https://logseq.com/
[8]https://tana.inc

that can be interlinked into knowledge and organizational structures; others, such as Tinderbox[9], implement visuo-spatial hypertext structures; some, such as Obsidian[10] and Logseq, emphasize privacy and data ownership, overlaying hypertext on plain text markdown files; and some, such as Notion[11] and Tana, natively emphasize combining linked notes with highly structured views like databases and timelines. These tools have substantial userbases, on the order of hundreds of thousands to a few million users, many of whom have been using these tools for a number of years.

We see a valuable opportunity here to observe and document real use patterns and impacts from usage of these tools, which may be meaningfully different from mainstream digital tools that are often the focus in empirical HCI research. These use patterns may also extend what was learned from usage of the earlier waves of hypertext tools like NoteCards. As far as we can tell, this opportunity has not yet been realized in the HCI community: some users of these tools have written essays to argue for the potential benefits of these tools for improving scientific work [16, 45], and these tools have begun to show up in passing mentions in (formative) empirical research of information management practices: for instance, one participant in [27] reported using Obsidian, though the details of their usage are not described.

We will discuss from our results how the open problems from the first waves of hypertext tools [25] remain or are addressed in the new wave of hypertext tools, and connect with other powerful design concepts in HCI, such as mixed-initative systems [28].

3 METHODS

To address our research questions, we conducted guided tour studies [21, 58] with 23 scholars, with a focus on how they use tools to augment their sensemaking work.

3.1 Positionality

All authors of this paper are active users of some of the hypertext tools studied in this paper. In addition, the second, third and fourth authors have participated in creating resources for their respective user communities, such as blog posts, Twitter threads, Youtube videos, and software extensions for the tools. This positioned us well to recruit suitable participants, and understand the nuances of the features and use patterns they share with us.

3.2 Participants

Our core sample was 16 participants who used new-generation hypertext notebooks for scholarly sensemaking. We purposively sampled from our direct contacts in the user communities of our respective tools (specifically Roam Research and Obsidian), as well as through social channels where the user communities gather, such as tool-specific Slacks and Discords (which we were able to access as active users), and on social media. We supplemented this core sample with a purposive sample of 7 participants who did scholarly sensemaking with "mainstream" tools, such as Google Docs. For analysis purposes, this supplementary sample includes one participant who used a hypertext notebook, but without using any of its hypertext features (e.g., not linking notes). Table 1 describes our core and supplementary samples in terms of the tools they used for their sensemaking work, and the general domain or research area of their scholarly work. For clarity, hypertext-using participants are prefixed with "H-", and non-hypertext users with "NH-".

3.3 Procedures

Guided tours were conducted remotely via video meetings. Participants first described the research directions for which they are conducting sensemaking work , which grounded the guided tour of their sensemaking practices. Then, they shared their screens and walked through their sensemaking workflows. Throughout, we asked follow-up and probing questions about changes to their workflows/systems, sources of structure and the forms of artifacts, and remaining pain points. Each interview lasted approximately 1-1.5 hours. All procedures were approved by our institution's IRB.

3.4 Analysis

All guided tour recordings were transcribed in full. We then conducted thematic analysis [7] of the transcripts alongside the screen recordings. Specifically, the first author performed open coding on the transcripts and captured screenshots of participants' shared screens displaying their work processes, and iterated on the codes through discussion with the other authors. For example, early bottom-up open codes focused on lower-level use patterns like *linking* and *querying*. Later, through constant comparison [14], discussions, and memo-ing, we iteratively induced higher level themes, like the "*hub*" and "*incubator*" meta-use patterns, and core problems like "*spatial fragmentation*" and "*temporal fragmentation*".

4 CONTEXT: CORE HYPERTEXT PRIMITIVES

To contextualize our core results, we first share a brief primer of key primitives of the hypertext notebooks our participants used.

Note Pages and Blocks. Note pages are often the smallest unit of the hypertexttools in our study: they are often the primary "node" that can be hyperlinked and composed into larger hypertext structures[12]. Each note page can contain one or more "**blocks**". Blocks can be a line or paragraph of text, or a bullet point in an outline. Importantly, in some tools like Roam Research and Logseq, blocks can also act as a unit for linking. And in some tools like Tana, they are the only unit for linking: Tana has no concept of pages.

Links and Contextual Backlinks. Links connect related note pages or blocks together. A typical way to construct links is to write a 'wikilink", which is a string enclosed in double squared brackets (e.g., "[[Linked Note Page Title]]"). In some tools, users can also create links between notes using "Markdown links", which are single square bracket-enclosed text followed by parenthesis enclosed URLs (e.g., [Link Name](AddressOfLinkedFile.md)). Clicking on these links will bring up the targeted note page.

Importantly, in these hypertext tools, links are not just one-way. Instead, these tools surface "*backlinks*" for a given note/block, which are other notes/blocks that link *to* that focal note/block. Backlinks are often populated by a native query system that searches for

[9]http://www.eastgate.com/Tinderbox/
[10]https://obsidian.md/
[11]https://www.notion.so/

[12]In hypertext parlance, these correspond to the concept of a "lexia": https://www.technorhetoric.net/1.2/features/eyman/lexia.html

Table 1: Participants and tools in our sample. "H-" and "NH-" prefixes denote hypertext and non-hypertext users.

Participants	Occupation	Research Interest and Direction	Tools
H1	Independent Researcher	Cognition, Learning, and Creativity	Bear
H2	Research Fellow	Politics and Economics	Logseq
H3	Independent Researcher	History	Obsidian
H4	PhD Student	Embodied Cognition and Language Understanding	Mindomo, Notion
H5	Professor	Scientific Information and Digital Publishing	Zetllr, Cosma
H6	PhD Student	Computational Neuroscience	Tana
H7	Professor	Computational Biophysics	EMACS with Org-Roam
H8	Professor	Cell Biology	Roam Research
H9	Research Staff	Cell Biology	Roam Research
H10	PhD Student	Organizational Behavior	Roam Research, Tana, LiquidText
H11	PhD Student	Older Adults and Caregivers	Logseq
H12	PhD Student	Interaction between Language and Cognition	Tana
H13	Professor	Film Studies	Roam Research, Freeplane
H14	Professor	Digitally mediated Music Engagement and Learning	Obsidian
H15	PhD Student	Computational Neuroscience	Notion
H16	Post-Doc Researcher	Health Research	Roam Research, Logseq
NH1	Independent Researcher	Expertise and Management	Ulysses
NH2	Researcher, Operator	Biology and Longevity	Blogs,Notion
NH3	Professor	Innovation and Economics	Blogs
NH4	Professor	Crowdsourcing and Citizen Science	Printout, PDF Reader
NH5	Startup Founder	Applied Math and Psychology	Google Docs
NH6	PhD Student	Natural Language Processing	Google Docs, Google Sheets
NH7	PhD Student	Computational Creativity in Music	Obsidian

other note pages that contain an outgoing link to the current note page. These backlinks are usually displayed as part of a tool's user interface, separate from the content of a note page/block. Some tools like Roam Research display backlinks at the bottom of a note page; others like EMACs with Org-Roam display backlinks in a separate window. In some tools, backlinks are a hidden feature that requires extra effort to turn on. For example, Obsidian requires users to select the "Backlink in Document" option in its setting menu display backlinks at the bottom of a note page; otherwise, backlinks are hidden in a tab in the right side panel. Bear originally had no native backlink feature, so our participant H1 built his own Bear plugin to enable backlinks[13].

A common feature that distinguishes the backlinks from earlier generations of hypertext is that they are *contextual*: they display backlinks in the context of the surrounding text in which the link appears. In contrast, mediawikis (which power Wikipedia, one of the more common manifestations of hypertext) display backlinks as a list of pages that link to a given page (without surfacing the context of the mention). In block-based tools like Tana, Roam Research, and Logseq, users are also able to traverse additional context of backlinks through their parents in the outline hierarchy.

Structure Searches. Different from direct searches, structure searches retrieve past notes using attributes and structural features of a note (in addition to its content), such as metadata, tags, and (importantly) hypertext structure, such as parent/child relationships between blocks, or links between note pages/blocks. We describe this feature as structure searches after Halasz et al's [25] description of the need for such a feature, as distinct from simple "content-based search", such as full-text search over the *contents* of each note in the hypertext notebook. These structure searches can be constructed using the native query language of the underlying database (e.g., `datalog` for Roam Research and Logseq), native query-building features (as in Tana), or with software extensions that enable users to write native queries in a user-friendly interface, such as the QueryBuilder extension for Roam Research[14].

Node and Link Typing. By default, many tools do not enable users to add formal "types" to their nodes and links. As we will see later in §5.2, users enact a form of node typing using note title prefixes. A notable exception is Tana, which includes a *Supertag* feature. Unlike normal tags, applying a supertag to a Tana block turns it into an "object" [15], enabling the block to be associated with a set of structured fields according to the supertag's schema. These fields can be populated with single values, but also with *links* to other Tana blocks, as well as embedded *structure searches*. In this way, supertags enable users to create typed nodes and links natively. A similar functionality is present in Roam Research with the QueryBuilder extension, which allows the user to specify formal node types and link types to use in structure searches.

Canvas. Some hypertext tools enable users to add or embed notes in an infinite canvas and draw connections between the notes. Users can follow links to the original notes from the canvas, just as if

[13]Recent updates of Bear now incorporate backlinks as a native feature

[14]https://github.com/RoamJS/query-builder

[15]https://tana.inc/docs/supertags

they were linked to a regular note page. In this way, they are a form of spatial hypermedia, similar to HyperCard or gIBIS [13]. Canvases that are native features of the hypertext notebooks (as in Obsidian and Logseq) usually allow users to place (or link to) note pages on the canvas. Not all tools, however, recognize the connections on the canvas as formal links/backlinks: one exception is the canvas feature in Roam Research (e.g., Fig. 2), enabled by the QueryBuilder extension. In QueryBuilder's canvas, users can create formal typed links between nodes (according to a user-defined schema), and these links are recognized in Roam as typed (back)links.

Graph View. Graph views are visualized networks of note pages and their interconnections. Some tools enable users to display 2nd- or nth-order links (e.g., notes that are indirectly connected to the current note by 2 or more degrees of separation) Obsidian, Roam Research, and Logseq natively support displaying the connections between note pages using a graph view.

5 RESULTS

Our results are fourfold: we describe 1) the nature of our participants' sensemaking work and the core problem of **temporally and spatially fragmented information** (§5.1), 2) the temporal dynamics of our participants' sensemaking work and how they manage temporal fragmentation and develop ideas over time by creating, retrieving, and developing different **note types** (§5.2), 3) how participants created **associative hypertext structures** — hubs, indices, and incubators — to manage temporo-spatial fragmentation (§5.3), and 4) how these hypertext use patterns contrast with how non-hypertext users deal with fragmentation (§5.4).

5.1 The Core Problem of Temporo-Spatial Fragmentation of Sensemaking Work

As we might expect, participants shared different ways that they do active, in-the-moment sensemaking, using work patterns reminiscent of prior work, like spatial canvases, and comparing things on the screen across multiple tools. H8's workflow for active sensemaking illustrates this pattern well: his lab uses a shared lab notebook in Roam Research to record evidence and claims from papers alongside results and conclusions from experiments. Figure 2 shows a canvas used to make sense of a biological model alongside associated questions, hypotheses, and evidence. All cards are linked back to note pages in Roam Research, which enabled him to draw connections between cards, make sense of information from the paper, and relate claims and evidence to their research questions and hypotheses in the shared lab notebook.

However, the majority of our guided tours were spent discussing "meta-work" required to get into a position to do this active sensemaking. In particular, our analysis of the use patterns the participants described clustered around a common problem of revisiting and reusing ideas, either before or during an active sensemaking session. This core problem was expressed in terms of ***temporo-spatial* fragmentation of sensemaking work**.

5.1.1 Temporal Fragmentation of Sensmaking. Participants described how sensemaking work naturally fragmented over time due to variations in available information, time, and energy. For instance, H1 described how he created a short question-style note to record his

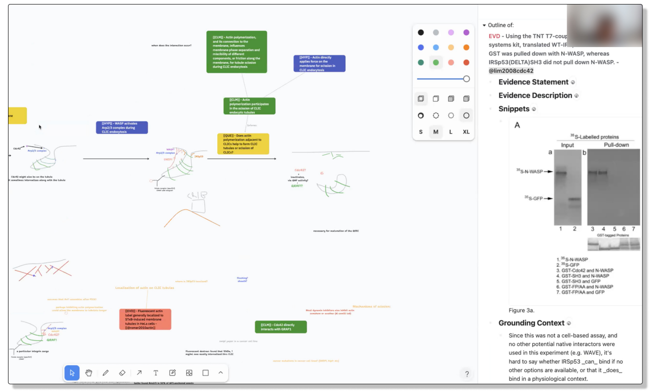

Figure 2: H8's active sensemaking of a biological model and associated claims (green), questions (yellow), hypotheses (blue), and evidence (red) on an embedded canvas in Roam Research. All colored cards link back to actual note pages (shown here on the right panel).

uncertainty about the evidence in a paper he was reading, as a form of *"just-in-time note-taking"*. He emphasized that this short note would not be the end of note-taking, elaborating:

> *"...if on a given day, I'm writing about retrieval practice and transfer learning, or I'm consulting that [question style] note, or I'm finding myself curious about that [question]... but not in a very predictable timeframe. This terrain pile will kind of get churned through. And the rest of the material will get incorporated."*

Many other participants, though, experienced temporal fragmentation due to distractible environments and interruptions, and often experienced the fragmentation as a drag on their sensemaking work. For instance, H2 described how his sensemaking was often interrupted by other concurrent responsibilities he faced from his job: *"I'm writing but I'm also managing people but I'm also in the organizational strategy stuff, and I'm hosting one or two podcasts a week and so it's it's a lot of things dragging my attention."* This distractible work environment impeded his ability to maintain a stable line of thinking on large projects:

> *"I found what I need for a big project: I need momentum. But that momentum, you know, through things outside of my control gets derailed pretty quickly. And then by the time I get back to it, I have like changed my mind about the whole project and how I want to be approaching it."*

H14, an associate professor of music, echoed a similar sentiment of how time got "scattered" during normal semesters. And NH3 similarly described how misalignments in sensemaking context between the current and the past required him to go *"back to the original source because the context noted down when reading the source does not match the final use context."*

This temporal fragmentation was closely tied to an expressed sense of fragility of relying on memory for ensuring retrieval of relevant information artifacts for active sensemaking, as exemplified by our non-hypertext participants. For instance, NH7 had a less distractible environment primarily focused on one research project, and did not create or use links to tie note pages across large projects together, instead relying on memory to return to what was most important. This approach, while sufficient for present purposes, was tinged with a background fear of forgetting: "*"Am I scared of forgetting things? Yes, I think, probably. But is the fear of not remembering greater enough than my inhibition to actually put in the work? No, not yet...that threshold hasn't been crossed."*. NH3 similarly noted, *"it might be hard for me to rediscover [interesting things I've seen] in two years"*. This fear of future forgetting was clearly underscored in NH1's recollection of "*deep regret*" for not making annotations on a book that was important for his research.

5.1.2 Spatial Fragmentation of Information Artifacts. Participants also described a tension between the required centralization of information artifacts for synthesis and the reality of spatial fragmentation of these information artifacts across a range of "information scraps" [6] and tools.

Information could be spatially fragmented between a primary note-taking tool and a mobile note-taking tool. For instance, H3's primary note-taking tool was Obsidian, but she would use Otter.ai to create audio memos of "literally just me talking to myself in the car" while picking up her son from daycare. Similarly, NH5 primarily used Google Docs to record notes for writing, but strayed from this when on the go: *"whenever I have a thought about it, I just have a place in my phone, a notepad, and put it in there."* Even within a single tool, information was often fragmented such that users frequently needed to do active search and retrieval work to find relevant notes. For example, H1 told us, *"I kind of go around my system and look for some other places to link to this (note page), because it's come up in a number of places."* Similarly, H3 needed to go over all her notes *"in case I came across something else that is relevant...That happens all the time."*

One cause of this fragmentation was a lack of organizational structure. This was often expressed by our non-hypertext participants, who showed multiple places for keeping notes, often describing a sense of messiness or chaos in their personal information management system. For instance, NH6 stored *"think things"* in a very *"disorganized and freeform way."* This spatial fragmentation led to frequent mentions of the difficulty of retrieving relevant information artifacts. In a previous iteration of his workflow, NH3 maintained a single note page in Roam Research to centralize all reading materials he found interesting on the Internet. Unfortunately, this note page gathered too many reading notes and became too long, making reading materials spatially fragmented in the vertical direction of this note page. During the interview, NH3 tried to find the one that appropriately triggered his memory when looking at a reading list in his abandoned Roam Research notebook, and complained, *"They get buried. I'll never find them again."*

But spatial fragmentation could also be a result of the temporal fragmentation of sensemaking. For instance, H14 kept "To Sort" folders in DevonThink to temporarily hold newly added PDF document, with an intention to later move these PDF documents to a more specific categorical folder; he later admitted, *"you could tell that's a lie because I have stuff going actually 8 years ago (in this folder)."* Similar to H1, NH5 had these "just-in-time" notes when he had an idea in mind and iterated them into more coherent essays over time. These notes *"capture[d] a lot of stray thoughts over a long period on the topic"*. However, this "just-in-time" note-taking created hundreds of draft notes in his Google Drive. He said, *"as you can see, I have like 130 unreleased finishes, that's just not great."*

As we will describe next, our participants developed a range of use patterns — leveraging key hypertext primitives like backlinks and structure searches — to overcome this temporo-spatial fragmentation of information.

5.2 Managing Sensemaking over Time with Explicit Note Types

One recurring use pattern amongst our hypertext participants was the creation and management of **explicit note types**. This use pattern had considerable range in its expression, from simpler, manual approaches like naming conventions and metadata, to more sophisticated, semi-automated approaches like note schemas.

A core property that interested us was how participants' note type distinctions mapped to key stages of development in the sensemaking process. As various models of sensemaking have outlined, sensemaking often progresses in loops of foraging and sensemaking over information artifacts that increase in their level of *structure* at the cost of increasing levels of *effort*. For instance, in Pirolli et al's [44] Notational Model of Sensemaking, intelligence analysts focus initial foraging efforts on searching and filtering external data sources into a *"shoebox"* of potentially relevant data, from which they read and extract information into an *evidence file*; later sensemaking activities schematize this evidence into a set of *schemas* and *hypotheses* that is eventually woven into a coherent *presentation* of a resulting case or story. Similarly, Qian et al [46] described how scholars would forage over initial *in-source annotations* on scholarly papers, from which they would develop *per-source summaries*, and eventually weave these into *cross-source synthesis notes* through sensemaking and synthesis work.

We similarly observed our participants creating explicit type distinctions between in-source annotations, per-source summary notes, and synthesis notes. For simplicity and continuity with relevant previous models, we describe this as a progression from external sources to **source notes**, to **summary notes** and **synthesis notes**, all of which are explicitly typed and integrated into hypertext workflows. As we will see later, enacting these note types within a hypertext notebook not only enabled our participants to track and manage the state of their thinking as it fragmented over time; it also unlocked powerful workflows for retrieving spatially fragmented information artifacts for active sensemaking.

5.2.1 From Source Materials to Summary Notes. A key early step in our participants' sensemaking workflows was the distillation of information from source materials, such as research papers, into **summary notes**: notes and comments on source materials.

A key distinguishing property — in contrast to marginalia in non-hypertext workflows — was that these summary notes were separate from the original source material; instead, these manifested as specific types of notes in participants' hypertext notebooks. Participants had different names for these notes, calling them literature notes, lecture notes, reading notes, and finding notes. Futher, in some workflows, participants would create **source notes** to represent source materials as reference-able, link-able objects in their

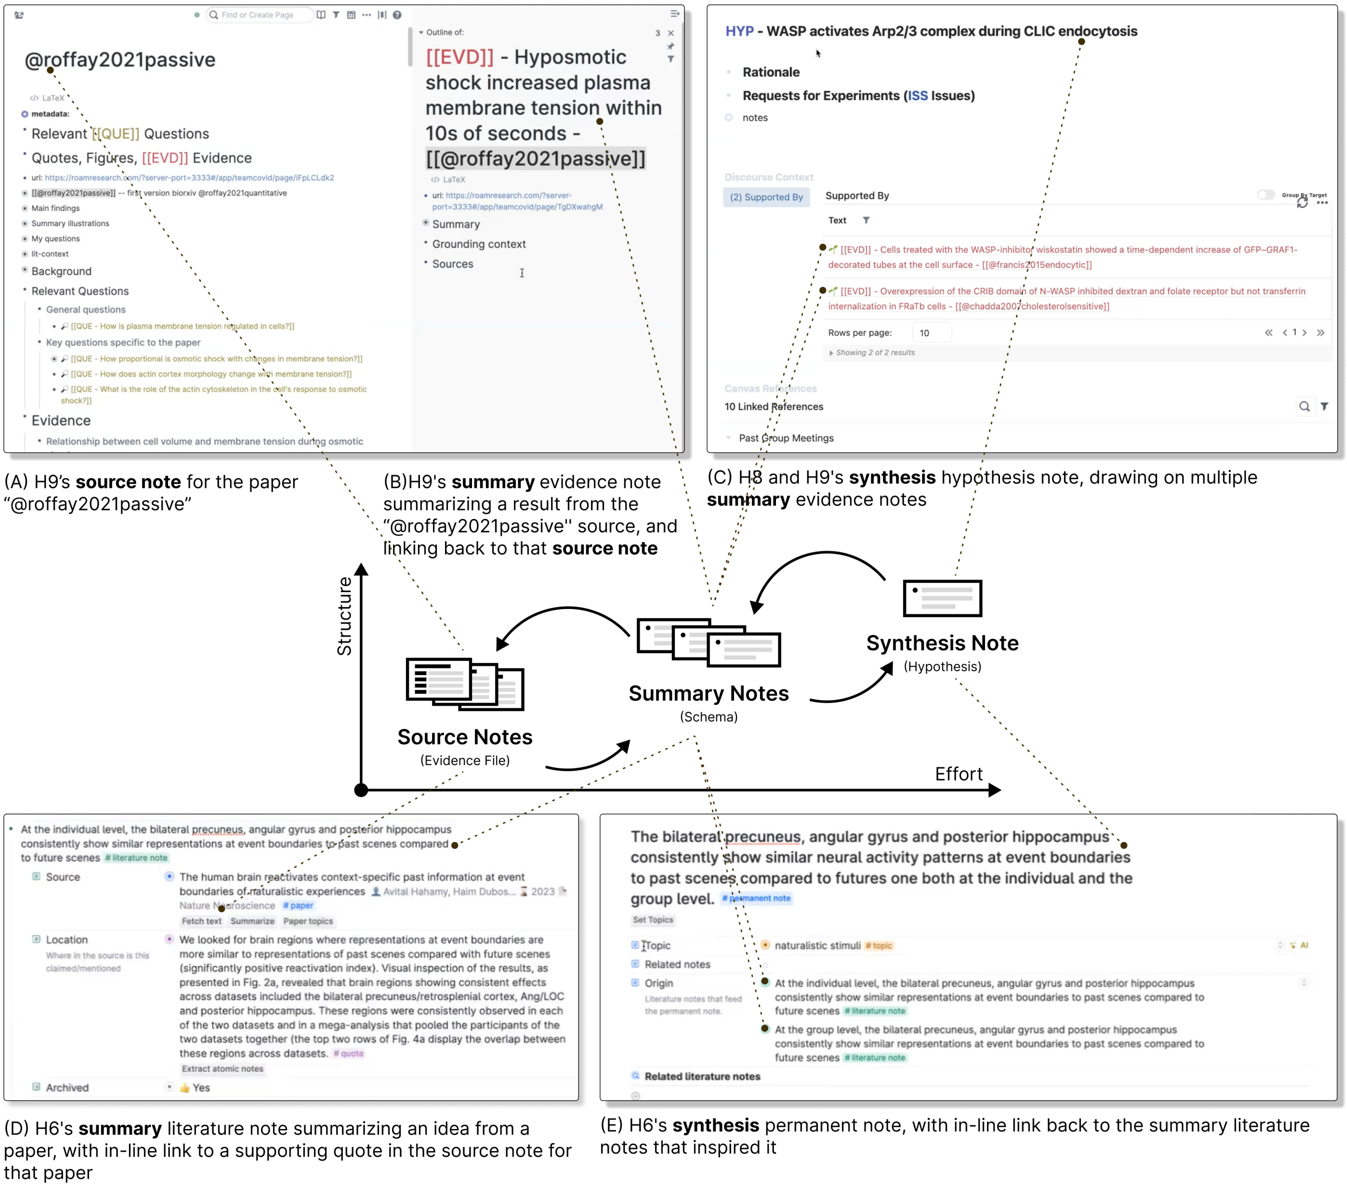

Figure 3: Canonical examples of explicit note types our participants enacted to match and progress the stages of their scholarly sensemaking, from source notes (A, D), to summary notes (B, D), to higher-level synthesis notes (C, E). Note how each later note type frequently included links back to note types from earlier stages.

hypertext notebook, which would provide more immediate access to provenance and metadata; other participants kept source materials separate from their notes, and were content with a textual pointer — such as a citation — to the source materials. All summary notes we observed, however, had the same core function of creating compressed information representations — with associated provenance tied to source material — that could then be used in sensemaking work within the hypertext notebook.

H8 and H9's "evidence notes" (summarizing a single empirical result from a paper) and "result notes" (summarizing a single empirical result from one of their lab's ongoing experiments) are canonical examples of summary notes. These note types were enacted using naming conventions — pre-pending *"[[EVD]] -"* to evidence note titles, and *"[[RES]] -"* to result note titles. By convention, evidence and result note titles also appended a hyperlink to the source note pages, encoding a preview of the provenance of those summary notes, and more details in the body of the source notes accessible one click away, if needed. Figure 3B shows one example evidence note: *"[[EVD]] - Hyposmotic shock increased plasma membrane tension within 10s of seconds - [[@roffy2021passive]]"* can be traced to the source note for the 'roffy2021passive' paper (Fig. 3A).

H6's "literature notes" are a closely related type of summary note: each note is an atomic summary of a key idea from a research paper (see Fig. 3 D). H5's "lecture notes" also summarize a single idea as described in a particular source, often with quotations from the text. In both cases, users enacted note typing with special metadata rather than naming conventions: in H5's case, a note type field, and in H6's case, a Tana "supertag" (which applied a "#literature note" type schema to his literature notes).

Some participants used semi-automated workflows for efficient, standardized note typing: for instance, H5 would create his notes in a command-line utility that would auto-fill metadata, including the note type; for H6, adding the "#literature note" supertag marked as a literature note type, with associated metadata and fields; H8 and H9 used a software extension that allowed them to mark snippets of text with a note type by highlighting the text and triggering a shortcut which would then create a new summary note with the appropriate prefix and metadata, as appropriate.

5.2.2 Consolidating Summary Notes into Synthesis Notes. Later stages of sensemaking for our participants were marked by the creation and development of **synthesis notes** that expressed higher level insights, conclusions, concepts, or hypotheses. These were often constructed from, or in response to, one or more summary notes; and sometimes woven together from other synthesis notes. As with summary notes, synthesis notes had different names and manifestations in participants' workflows, such as topic notes, concept notes, question-style notes, project notes, and write notes. Similarly, synthesis notes were typically marked explicitly as such, using similar mechanics of naming conventions or metadata/schema.

One example of forming synthesis notes from summary notes can be seen in H6's adaptation of a networked note-taking approach for sensemaking called the *Zettelkasten*, or "slip box" [2, 22]: he would use one or multiple literature notes to form a permanent note (Figure 3E), which he said were atomic notes — often expressing his own original ideas, or articulations of important ideas in his own words — that are ready to use in his writings in the future. H6 found it important to formally distinguish between these note types by applying a Tana "Supertag" for either "#literature note" and "#permanent note". As we will see in §5.3.2, this formal typing enabled H6 to construct powerful views of his notes using structure searches that leveraged this note type metadata.

After a literature note was used to form a permanent note, H6 would then mark the literature notes as archived in the Archived field of "#literature note" Tana Supertag. By doing so, only unprocessed literature notes will appear in his topical index alongside permanent notes (Fig. 5A). These permanent notes and unprocessed literature notes would become materials for his writings. H5's workflow included a similar archiving dynamic for some summary notes; he described how *"eventually what [he'd] like to do is dissolve them into concepts, ideas, insights, everything. And maybe don't even have those lecture notes again."*

Hypothesis notes are another example of synthesis notes from H8 and H9's note-taking system, instantiating what they learned from source materials and their previous knowledge (see example in Fig. 3C). H8 created hypothesis notes after going across different source materials, saying *"if there isn't a satisfactory claim, we'll label the claim as a hypothesis. But it's sort of our best guess."* Instead of directly combining information from different sources, the formation of hypothesis notes marks the synthesis of what is unknown. These hypothesis notes would then ***"operationalize"*** their research questions. He elaborated, *"It's sort of directing us towards conducting some experiments that would give us the answer to the question or a more satisfactory answer to the question."* H8 said using the *"[[HYP]] -"* naming convention helped them *"have a good sense of how to test the hypothesis."*, calling these hypothesis notes *"actionable discourse nodes"*. H8 expressed that *"(in) absence of a system like (Roam Research), it's just too much information to try to process."*

H8 and H9's hypothesis notes demonstrate an interesting dynamic between synthesis notes and source materials: the production of synthesis notes from existing source materials pushed them to plan for the action of getting new source materials (i.e. designing experiments and collecting data). This dynamic encloses the material loop between source material, summary notes, and synthesis notes. This life cycle of notes was echoed by H1 (described in more detail in §5.3.3): *"It's usually inspired by some question that I have, both because that's kind of the new nucleation site for the synthesis, and also because it tends to be the generator for the reading queue."* This enclosed life cycle of note types depicts the material side of sensemaking. Source materials, summary notes, and synthesis notes are tied to — and indeed are the substance of — different stages of participants' sensemaking work. They are both a representation of and a response to the temporal evolution of sensemaking, and enable our participants to keep track of the various stages of their thinking over time. Indeed, the explicit note types signals the level of development of a given information artifact, such that when the user returns to it, they know what they can do with it, and what needs to happen next, if any (e.g., summary notes need to be consolidated to synthesis notes; synthesis notes can be directly used in active sensemaking on a canvas or in a draft to make outputs). But they also exemplify how discrete information artifacts in a hypertext-based note-taking system could be. In the next subsection, we will describe how participants utilized hypertext affordances to interact with these note types and manage transformations between types of notes, deal with temporo-spatial fragmentation, and push their scholarly sensemaking work forward.

5.3 Extending Associative Navigation with Hubs, Indices, and Incubators

Participants described using their tools' native link and backlink features to overcome the spatial fragmentation of their notes across different pages and folders within their note-taking system. By navigating through this associative structure of their hypertext notebook, participants were often able to retrieve information that was hard to search for. Beyond these simple associative structures, however, participants often constructed specialized information artifacts that extended their ability to associatively retrieve information that was fragmented across their hypertext notebook.

Here, we describe three recurring patterns:

(1) **Hubs**: empty note pages that use backlinks or queries to centralize distributed information artifacts.
(2) **Indices**: manually constructed note pages that contain links to a set of related notes, as well as descriptions of the nature of the information being linked to.

(A1) Backlinks to the hub note provide access to related important information, with context of relationships shown

(A2) The body of the hub note page is empty: semantics of the hub note defined by backlinks, not direct content

(A) H7's Empty Hub Note for A Concept

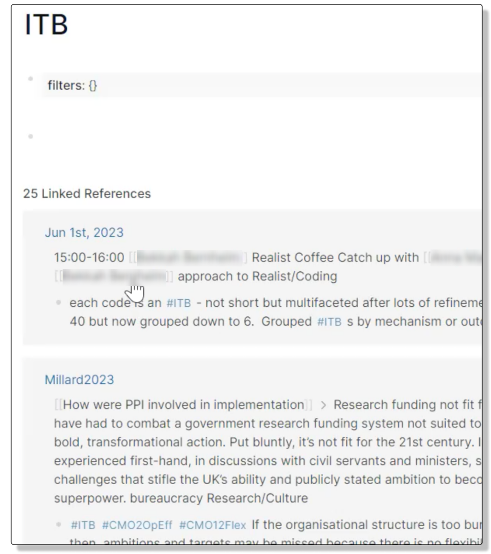

(B) H16's hub note is similarly empty, with context-rich backlinks to related information

Figure 4: Canonical examples of hub notes that are contentless, but whose backlinks provide context-rich waypoints to related information scattered in a hypertext notebook.

(3) **Incubators**: note pages where participants add and describe links to other notes, and also add a *temporal or developmental dimension*, such as distinguishing between stages of maturity of ideas, articulating open questions or tasks for further development, or maintaining the provenance of information on a timeline of activity.

5.3.1 Hubs. A hub is a note page that is linked to from multiple other note pages, providing a waypoint to other note pages via the hub's backlinks. Users typically create a hub note by creating a link to it in a target note, rather than creating it separately with its own content: the hypertext notebook then creates the page separately, and populates its backlinks section with a link from the current note. Thus, hubs typically contain no manually typed annotations or notes: instead, they are empty note pages that are constructed automatically through backlinks.

A representative example of hubs is H7's case of using backlinks on an empty note page. H7 was a professor at a university in computational biophysics, using software called EMACS with Org-Roam as his personal notebook. He created an empty page (Fig. 4A2) for the topic *"Agent-based models"* that elicits backlinks (Fig. 4A1), elaborating that *"(it) has no other purpose than being the hub for topics."* H7 would then reference this hub note in notes on articles he collects and reads, thereby creating links between that page and those articles. The *"Agent-based Models"* hub page would then be an access point to these other resources — such as a page on *"computational scientific models"* — via its backlinks. These backlinks provide more than lists of pages that links to the current hub note: instead, they show the rich *context* of the backlinks via the text that surrounds the link to the hub note. This context provides additional information about the nature of the link, as well as in-line access to related information. In this way, each of these backlinks can in turn can be traversed to other information artifacts, such as important authors (e.g., *"Thomas Schelling"* under the backlink for *"Dynamic models of segregation"*; Fig. 4A1) as well as other articles or topics of interest. Describing this process, H7 said:

> *"When I get some new item in my database...about agent-based models, I make sure to put a link to agent-based models. So there are empty pages, basically, in my wiki, having no other purpose than being the hub for topics."*

Most hubs follow a similar pattern: for example, H5 and H16 also create empty note pages to maintain connections (H16's example shown in Fig. 4B); and H8 and H9 create empty project pages that collect information about a project via its backlinks.

5.3.2 Indices. Participants also created indices, which centralized links to other notes in a more structured and descriptive manner, such as annotated links in the body of the index note page, or structure searches displayed in the form of lists and tables; this intentional, active structuring contrasts with the more passive, structureless nature of hubs, that provide a flat list of waypoints to other note pages via backlinks.

Indices can be as simple as elaborating a hub note manually with links in the body of the note page. For instance, H10 manually created an index of topics on an index page, ordered in alphabetical order, to *"gather together all of the topics that interest you until you can refine them further."* (see Fig. 5C).

More commonly, our participants would construct and maintain indices in a semi-automated manner by embedding *structure searches.* The structure searches extend the functionality of the native backlink queries, which are essentially a flat query for links

(A) H6's semi-automated index of permanent synthesis notes and summary literature notes populated by structure search for each topic

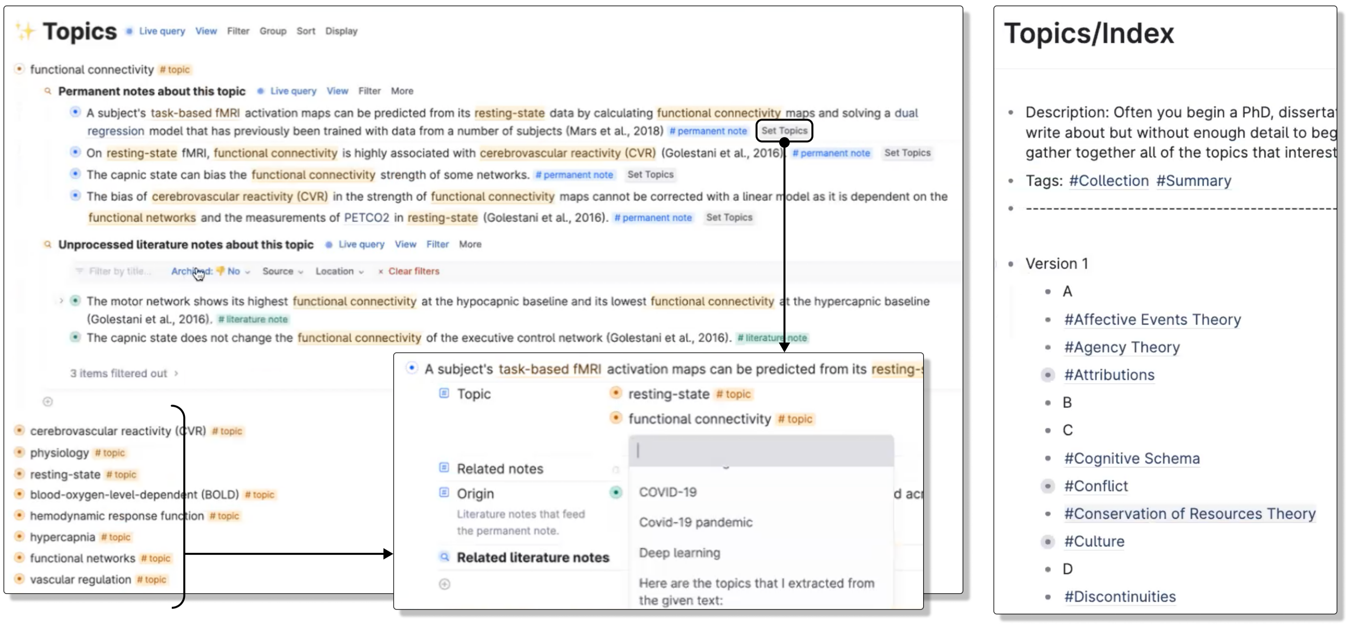

(B) H6 can populate the topic field for each permanent/literature note manually or with AI-assisted "Set Topics" command that draws from topics already in the index where possible

(C) H10's Manual Index of Topical Tags

Figure 5: Two canonical examples of indices created by our participants, using sophisticated semi-automated (A, B) and simpler manual approaches (C).

to the note; structure searches, by comparison, enable more sophisticated sorting and filtering by different semantic categories or even types of relationships. One striking example of such semi-automated indices is H6's topical index. H6 was an advanced PhD student in Computer Science and Neuroscience. He created an index of topics in Tana to manage literature notes and permanent notes that are core to his sensemaking process (see Fig. 5A). This list of topics gathered spatially distributed literature notes and permanent notes. As we described in Section 5.2.1, literature notes were texts that summarize and comment on the quotes he captured from papers and articles, tagged with the *"#literature note"* Tana Supertag. Permanent notes were texts that consolidate ideas from literature notes, followed by *"#permanent note"* Tana Supertags. These permanent notes would be directly reused in his writings later. H6 would use this topical index to filter notes he took in the past by topics on this note page and have everything ready at hand for active sensemaking, such as when writing the introduction or related work for a grant or research paper: *"I would open this (topical index). It would be like, "Oh, this sentence I would like to use in this manuscript" and then I would copy and paste.""*

H6 had a sophisticated mixed-initiative approach to maintaining this topical index. The index itself was a structure search of all "#topic" notes; indented under each topic were structure searchs for literature and permanent notes whose topic field included that topic. While H6 could manually populate these topic fields when creating literature or permanent notes, he would more commonly use an AI-assisted approach: his *"#literature note"* and *"#permanent note"* supertag structures included a "Set Topic" button, which, when clicked, would trigger an OpenAI routine to automatically populate the Topics field of that literature or permanent note with a topic note based on the content of the permanent note (see Fig. 5B). Importantly, H6 said, *"I have set up the command for OpenAI...to, if possible, use a list of topics I already have."* This enabled him to control the complexity and consistency of his topic index.

H6 also used Tana Supertags to structurally query notes for creating an idea index and paper index, so he could view past ideas and track his reading process. Other participants (H8, H9, H11) similarly used structure searches to create indices that monitored where papers were mentioned, gathered specific note types, and dynamically retrieved open TODOs.

5.3.3 Incubators. Our participants also created incubators: note pages where they more actively added information and links to other note pages over time. Compared with hubs and indices, incubators include richer information than associative connections, reflecting participants' most up-to-date thoughts on a project, a topic, a concept, or a question. Participants described how they made notes and comments in the incubators over time. Incubators were therefore critical constructs for dealing not just with spatial, but also temporal fragmentation, of sensemaking work.

One salient example of incubators is H1's log-type notes for recording questions and monitoring experiments. Log-type notes are a type of note page that records notes in chronological order. H1's log-type notes would include entries organized by date over

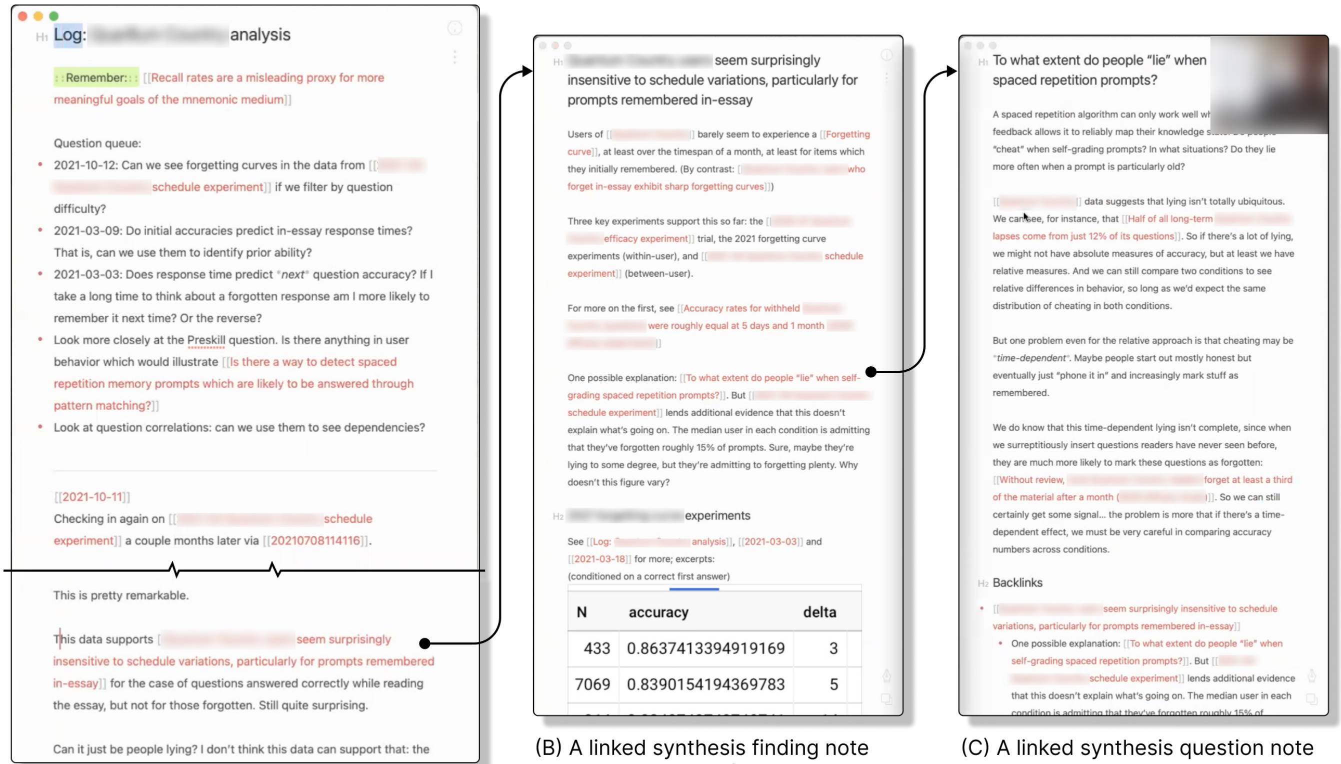

(A) H1 recorded thoughts about a research project over time in a log-type **incubator** note

(B) A linked synthesis finding note synthesizes results from the project

(C) A linked synthesis question note synthesizes alternative explanations for the finding note

Figure 6: An example of a manually-created, log-type incubator note from H1, linking to and discussing links to other information over time to make sense of ongoing research questions and results.

time (see Fig. 6) and were the go-to place for H1 to keep track of his thoughts: he described this type of note as *"something that I come back to again, and again, asking for help, monitoring experiments and whatever."* Describing his log-type incubators' rich contextual structure and dense links to other notes, H1 said:

> *"It's connected to a bunch of these findings and questions...[and] these so-what headlines will bubble up into things like this. And there's some even higher level things in (it)."*

More interestingly, the hypertext links in H1's incubators are not in the form of short noun phrases as in Wikipedia, which is only several words long, but in the form of complete sentences that summarize the key idea of the linked page in a succinct declarative statement. Using other plain text that surrounds the hypertext links, H1 provided even more details and context of why he chose to link in a particular note page. For example, the surrounding text for the inline link to a synthesis finding note (Fig. 6B describes how the linked finding note provides further support for the current observation; an inline link from that finding note to another question note (Fig. 6C), along with its surrounding text, elaborate on an alternative explanation for the findings.

H1 further described how the interlinked note pages of findings and questions revolving around the log-type notes enabled him to manage the temporo-spatial fragmentation of his sensemaking:

> *"The idea is that synthesis is happening here. It's over time rather than over space – across papers where thousands of words in here are bubbling up into these kinds of medium-form things which are getting connected to some papers sometimes… And eventually, this will become some large publication."*

Highlighting how this incubator log-type note helps address temporal fragmentation, H1 further elaborated, *"I haven't updated this (log-type note) in a year or so. But it still applies pretty well."*

Incubators in H1's case took advantage of hypertext and plain text, weaving a complete record of his thought of thinking. As we saw in this example, log-type notes became the center for H1 to go back to temporally fragmented sensemaking work and gather spatially separated information artifacts, pushing H1's scholarly work forward. In H1's words, the log-type notes are *"usually inspired by some question that I have, both because that's kind of the new nucleation site for the synthesis, and, also because it tends to be the generator for the reading queue."*

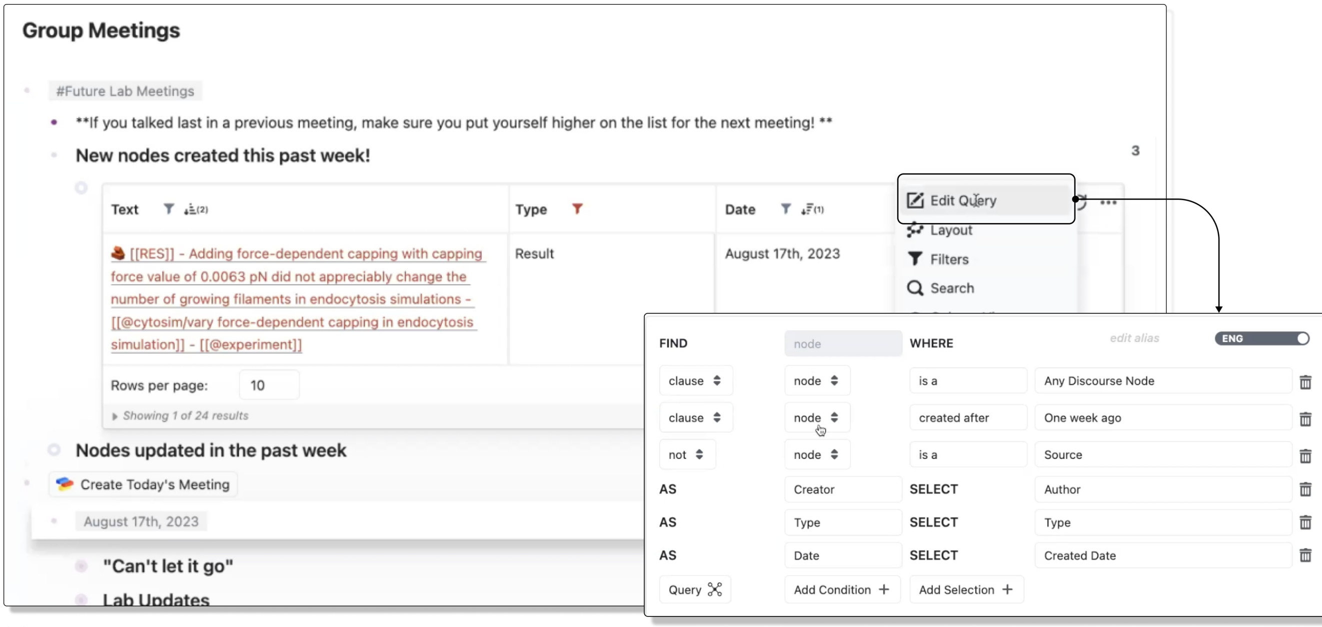

(A) H8 and H9's **incubator** group meeting notes page

(B) The structure search for pulling in information into the incubator can be customized to satisfy different information needs

Figure 7: An example of a semi-automatically constructed incubator note: a group meeting notes page that incorporated a structure search to pull in related information, but also including manually curated links and information over time.

Similar to H1's largely manual, active construction and development of his incubator, many other participants manually maintained incubators that integrate content about key ideas, quotes from source materials, and participants' comments on these ideas and quotes. For instance, H10 — in contrast to her more passive and stable indices — incrementally built her project pages, course pages, and definition pages over time in Tana, by building links with source pages at different levels of abstraction, and even outlining some thoughts in the incubator. In her own words:

> *"I'm not interacting with a direct source. But I'll start to outline certain things there. And then I'll actually differentiate, so the next thing I like about this is that I have like a history of workplace conflict that's, I'd say, very abstract. It's at a more abstract level. And then, I get a bit more space. So here are different definitions."*

H14 also kept verbatim quotes and key ideas on a concept page and used links to connect these quotes and ideas to his source pages on an infinite canvas page in Obsidian. H13 used topic pages and reading source pages in Roam Research to keep ideas, comments, and quotes. H15 used Notion to actively gather, organize, and elaborate on links to — and quotes from — note pages about a range of topics and resources relevant for an ongoing concept she was researching for her PhD thesis. And H9 maintained experiment pages as incubators, not only logging details of her ongoing experiments in a chronological structure, but also linking to and elaborating on relevant pictures from a physical lab notebook, images from experiments, and related question/hypothesis/evidence pages. This cluster of related cases illustrates the range of variation in types of structures within incubators, from chronological or developmental progressions, to levels of abstraction, to types of resources.

As with indices, however, incubators could also be constructed and developed in a semi-automated fashion. H8 and H9's group meeting note page is an example of this pattern (Figure 7). On top of the group note page is a table view of query results that H8 could filter through to retrieve relevant past group notes for writing. Different from indices, the group meeting page captures information created by different members of the lab over time, collectively incubating ideas for the entire lab. The key element of this page was the record of meeting notes (Fig. 7A). The query shown in Figure 7B was just a part of the incubator that provided H8 and his team a quick view of what they had produced over the past week. This example highlights that an incubator can be composite. Features from hubs and indices could be seen in incubators. The incubator can even evolve from indices and hubs.

5.4 Managing Temporo-Spatial Fragmentation without Hypertext

Our identification of hypertext hubs, indices, and incubators sensitized us to notice that our non-hypertext participants also created artifacts similar to indices and incubators (but not hubs).

For example, some participants, like NH2, constructed tables to make indices of the papers they collected, recording metadata like the name, URL, summary, publication year, and local context like keyword tags and the date added. In Notion, these keyword tags could be used to filter out papers about specific tags: as NH2 elaborated, *"you're doing a filter by that, look at those papers, right,* *the section, the big intersection, and then complete section by section."* Some other participants created lists to index information artifacts. For example, NH7 created a reading list as a note page in Obsidian, which included hyperlinks to the PDF files of the papers, different from associative connections between different note pages in a hypertext notebook. NH3 created a list of article ideas for future writings in Microsoft Word, ordered by priority.

Participants also tracked the status of information artifacts without using hypertext, marking an alternative life cycle of information artifacts. This could be accomplished with local color-codes, labels, or tags. But some strategies also increased spatial fragmentation: for instance, NH6 tracked the status of her reading notes in part based on where they were stored: in a Google Sheet, an annotated bibliography of a selective set of readings in a Google Doc (organized by the order of her planned writing work), and in Overleaf documents. Expressing the spatial fragmentation of this strategy, NH6 told us, *"It's just a lot of information to keep track of as I'm synthesizing. But I also needed to know all of the details. So, trying to look at this (Google Doc) document and then think about what's similar between these papers is really difficult. Like disorienting kind of. So then going from these (copying annotated bibliography to Overleaf)."*

6 DISCUSSION

In this paper, we explored the potential of augmenting sensemaking with hypertext, by analyzing in-depth guided tours of scholars who use new hypertext notebooks for their scholarly sensemaking (and comparing their experiences and use patterns to scholars who use non-hypertext tools such as Google Docs). We were particularly interested in comparing use patterns from this latest wave of hypertext notebooks to use patterns (along with the open issues identified) from the first generation of hypertext tools [20, 25, 59].

Our results reveal how the core problem of temporo-spatial fragmentation of information is a bottleneck to sensemaking, and a rich set of use patterns in hypertext that enable scholars to address this core problem. These use patterns — managing sensemaking over time with explicit note types, and extending associative navigation with hubs, indices, and incubators — reveal two key design patterns for hypertext-augmented sensemaking:

(1) empower users to flexibly define explicit types of information artifacts that match — and incorporate rich affordances for progressing forward through — the stages of their sensemaking process, and
(2) empower users to construct and use dynamic associative navigation structures.

6.1 Implications for Hypertext Support for Sensemaking

6.1.1 Extending known hypertext use patterns. Some of our identified use patterns echo what has been observed in earlier waves of hypertext use patterns, but with added nuances.

For example, hubs, indices, and incubators as associative navigation structures are reminiscent of overview-type notes from the previous generation of tools, such as Browser or Filebox cards in NoteCards [25], and Table of Content notes in Proteus [20, 59]. But there are some important nuances. For instance, the backlinks are more granular and contextual, in part due to the often more declarative note titles (e.g., for summary and synthesis notes), and the block-based structure of tools like Tana and Roam Research that expose the outline hierarchy position of note references in the backlinks. The indices and incubators we documented also have more flexible structures compared to the strict hierarchical organization in previous hypertext tools. In our participants' workflows, notes can have many-to-many relationships, including multiple "parent" notes. The incubators, too, jointly address both temporal and spatial fragmentation by providing an integrated way to retrieve information across the hypertext notebook, while also managing state (e.g., by logging progress over time, or distinguishing between materials that are of different stages of development). Finally, we observed participants leveraging powerful *structure searches* in their indices and incubators, and integrated them with AI-powered workflows (e.g., H6's AI-powered topic tagging workflows that drew from and contributed to his central topical index).

Separately, our participants' flexible specification of note types is reminiscent of the ability to specify new node types in previous systems. Here, too, we observed an interesting extension in the use of note schemas (with embedded, editable fields that could themselves incorporate dynamic structure searches), as seen, for example, in H6's use of "supertags" for note types.

6.1.2 Advancing open problems in hypermedia. A closer look at these use patterns and nuances in light of the key open problems for hypermedia discussed in Halasz's further clarifies how the current generation of hypertext tools have advanced the design space of hypertext-augmented sensemaking [25].

For instance, the inclusion of structure searches in the indices and incubators suggests that the new generation of hypertext tools has begun to address key open problems around searching and querying in hypermedia networks. The rich internal structures of the incubators — particularly the formally enacted supertag structures — also address Halasz's proposal to augment the basic node and link model with composite structures that have more sophisticated internal semantics. Finally, the structure-search-powered indices and incubators showcase progress towards dynamic "virtual structures" for dealing with changing information: as H6's dynamic index illustrates, the current generation of hypertext tools enable not just static queries, but dynamic views with sophisticated internal semantics that can be directly edited and manipulated. H6's index also illustrates some progress towards more sophisticated computation in (and over) hypermedia networks.

6.1.3 Future directions for hypertext-augmented sensemaking. Reflecting on our use patterns in light of old open problems also highlight future directions for hypertext-augmented sensemaking.

For example, we see a need for further work towards Halasz's [25] vision for tooling and practices that maximize extensibility and tailorability of hypertext systems. A substantial portion of our participants still largely used manual approaches to creating/managing note types or creating/using associative navigation structures. While the structures were often quite lightweight and interleaved with informal text, the reluctance of some of our participants to leverage these hypertext structures even when they were using hypertext notebooks (in favor of even faster and more lightweight reliance on memory, as with NH7), suggests that the cost of specifying even these lightweight structures may not be favorable in some cases. This problem recalls previously documented tradeoffs between the downstream benefits of organization and structure, and the need for flexibility and speed in the moment of sensemaking [32, 49]. Further, the query languages in the tools we observed — along with the programming-like experience of using them — may still be out of reach of most users. We see great potential in connecting this stream of research to work on end-user programming [33], to lower the cost structure of creating these powerful hypertext use patterns. H6's integration of AI assistance for adding structure to his notes also points the way to future research directions around adding structure in hindsight, leveraging the rich structural semantics of hypertext [37] as a *shared representation* [28] for human-AI interactive workflows.

Building on this, our participants' successful integration of AI assistance in their workflows — as seen in H6's AI-assisted topical index — further inspires us to envision future *mixed-initiative* [30] systems for hypertext-augmented sensemaking. For instance, our participants' summary notes and synthesis notes had a clear sense of progression in the sensemaking process, with specific affordances and actions for each note type that enabled forward progression (e.g., comparing and connecting summary notes into synthesis notes). Hubs, indices, and incubators, too, could be seen as a forward progression of more sophisticated and powerful associative navigation structures. Could we imagine note types that behaved more like "agents" with (pro)active behaviors to facilitate these forward progressions? Could claim or hypothesis notes actively scan for supporting/opposing evidence, both within the hypertext network of notes, but also in external sources, and draft associated note types and links to further develop the claim?

6.2 Limitations and Future Work

6.2.1 The cost structure of sensemaking with hypertext. While our study design enabled detailed empirical descriptions of use patterns (along with users' experiences and feelings around those use patterns), we do not have direct evidence of efficacy. As such, our results are vulnerable to the critique that the use patterns we identified may not actually be helpful for sensemaking.

We think this is unlikely for two reasons. First, the majority of our participants described their workflows as matter-of-fact parts of their main workflows, that were "working" for them; indeed, many also noted how their hypertext structures were crucial for managing temporo-spatial fragmentation and accomplish sensemaking work they were proud of. As a striking example, H2 shared how he *"built a comprehensive outline over time"* using Logseq, and put himself into an optimal position to take advantage of the limited time he was able to set aside for active sensemaking: "*I just was able to like crank out because I had this outline. So then it was just a matter of, you know, turning an outline into prose*."

With that said, it would still be valuable to conduct comparative experimental studies in the future to estimate or quantify the gains relative to the costs of implementing these hypertext patterns for sensemaking. Our description of the core problem of temporo-spatial fragmentation could help ensure such studies are conducted in the conditions under which we expect these hypertext use patterns to be useful.

6.2.2 Open problems around spatial fragmentation of information. Our results here also focus on how spatial fragmentation is dealt with *within* hypertext notebooks. The problem of spatial fragmentation *across* tools and devices remained unsolved for many of our participants: there is still a need for tools that can weave together data across the various tools and devices where relevant information may be stored. Still, we believe that the design pattern of incubators may be a useful structuring metaphor, regardless of the technical details of how linking across devices/tools/media is accomplished. We can see, for example, approaches for surfacing logical dependencies between files across tools [24], or AR/VR-based approaches to materializing information, being linked together not just in a static "master notebook" (as observed in [55]), but in dynamic incubators that can grow over time. In this way, the ideas behind hypertext could complement efforts to address cross-tool fragmentation in terms of data interoperability or mixed reality.

6.2.3 Extending beyond hypertext and scholarly sensemaking. Our findings are most usefully interpreted in the setting of scholarly sensemaking with hypertext tools. However, some of the patterns may have broader implications. For instance, participants' note types that mapped to stages of their sensemaking (§5.2) resemble phases of ethnographic practice [18], shifting from "jottings" (just-in-time, rough thoughts recorded in the moment, without regard for grammar or spelling), to "field notes" (summary or curated notes, with proper grammar and sentences) to "memos" or thematic analysis (synthesis or "theoretical codes", as in grounded theory [11]). The transition from summary notes to synthesis notes (§5.2.2) is particularly reminiscent of the processes of grounded theory [11] and the KJ method of affinity diagramming [50] in qualitative research. On the flip side, some work has also documented how scholars adapt qualitative data analysis software, like NViVo and Atlas.ti, for scholarly sensemaking [38, 61]. Thus, future work could explore how design patterns from qualitative analysis software and hypertext notetaking tools might usefully inform each other.

7 CONCLUSION

In this paper, we explored the question of how we might augment sensemaking with hypertext through an in-depth analysis of guided tours of scholars using new hypertext notebooks for their scholarly sensemaking. We identified novel design and usage patterns for dealing with the temporo-spatial fragmentation of sensemaking work, such as note types, and associative structures of hubs, indices, and incubators for flexibly retrieving information for sensemaking. These usage patterns — in part leveraging novel features of the new wave of hypertext notebooks, such as structure searches, and formal note schemas — expand the design space of approaches for hypertext-augmented sensemaking, and open up new research directions for the future of sensemaking.

ACKNOWLEDGMENTS

This research was supported in part by grants from the Graph Foundation and Sloan Foundation. We thank our participants for their insights, and Aniket Kittur, Max Kreminski, and our anonymous reviewers for helpful discussions that significantly sharpened the insights in this paper.

REFERENCES

[1] Mark S. Ackerman, Juri Dachtera, Volkmar Pipek, and Volker Wulf. 2013. Sharing Knowledge and Expertise: The CSCW View of Knowledge Management. *Computer Supported Cooperative Work (CSCW)* 22, 4-6 (Aug. 2013), 531–573. https://doi.org/10.1007/s10606-013-9192-8

[2] Sönke Ahrens. 2017. *How to Take Smart Notes: One Simple Technique to Boost Writing, Learning and Thinking – for Students, Academics and Nonfiction Book Writers.*

[3] Mark W. R. Anderson and David Millard. 2022. Hypertext's meta-history: Documenting in-conference citations, authors and keyword data, 1987-2021. In *Proceedings of the 33rd ACM Conference on Hypertext and Social Media (HT '22)*. Association for Computing Machinery, New York, NY, USA, 96–106. https://doi.org/10.1145/3511095.3531271

[4] Mark W. R. Anderson and David E. Millard. 2023. Seven Hypertexts. In *Proceedings of the 34th ACM Conference on Hypertext and Social Media (HT '23)*. Association for Computing Machinery, New York, NY, USA, 1–15. https://doi.org/10.1145/3603163.3609048

[5] Ofer Berman and Steve Whittaker. 2016. *The Science of Managing Our Digital Stuff* (1 edition ed.). The MIT Press, Cambridge, MA.

[6] Michael Bernstein, Max Van Kleek, David Karger, and M. C. Schraefel. 2008. Information Scraps: How and Why Information Eludes Our Personal Information Management Tools. *ACM Trans. Inf. Syst.* 26, 4 (Oct. 2008), 24:1–24:46. https://doi.org/10.1145/1402256.1402263

[7] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (Jan. 2006), 77–101. https://doi.org/10.1191/1478088706qp063oa

[8] Joseph Chee Chang, Nathan Hahn, and Aniket Kittur. 2016. Supporting Mobile Sensemaking Through Intentionally Uncertain Highlighting. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. ACM, New York, NY, USA, 61–68. https://doi.org/10.1145/2984511.2984538

[9] Joseph Chee Chang, Nathan Hahn, and Aniket Kittur. 2020. Mesh: Scaffolding Comparison Tables for Online Decision Making. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology (UIST '20)*. Association for Computing Machinery, New York, NY, USA, 391–405. https://doi.org/10.1145/3379337.3415865 00001.

[10] Joseph Chee Chang, Amy X. Zhang, Jonathan Bragg, Andrew Head, Kyle Lo, Doug Downey, and Daniel S. Weld. 2023. CiteSee: Augmenting Citations in Scientific Papers with Persistent and Personalized Historical Context. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–15. https://doi.org/10.1145/3544548.3580847

[11] Kathy Charmaz. 2014. *Constructing Grounded Theory.* SAGE. Google-Books-ID: v_GGAwAAQBAJ.

[12] Duen Horng Chau, Aniket Kittur, Jason I. Hong, and Christos Faloutsos. 2011. Apolo: Making Sense of Large Network Data by Combining Rich User Interaction and Machine Learning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 167–176. https://doi.org/10.1145/1978942.1978967 00000.

[13] Jeff Conklin and L Begeman. 1988. gIBIS: A Hypertext Tool for Exploratory Policy Discussion. *ACM Transactions on Office Information Systems* 6, 4 (1988), 29. https://doi.org/10.1145/58566.59297 00000.

[14] Juliet Corbin and Anselm Strauss. 2008. *Basics of qualitative research: Techniques and procedures for developing grounded theory (3rd ed.).* Sage Publications, Inc.

[15] Tim Coughlan and Peter Johnson. 2008. Idea management in creative lives. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems (CHI EA '08)*. Association for Computing Machinery, New York, NY, USA, 3081–3086. https://doi.org/10.1145/1358628.1358811

[16] Lance Cummings. 2023. Content Management System 3.0: Emerging Digital Writing Workspaces. In *Digital Writing Technologies in Higher Education : Theory, Research, and Practice*, Otto Kruse, Christian Rapp, Chris M. Anson, Kalliopi Benetos, Elena Cotos, Ann Devitt, and Antonette Shibani (Eds.). Springer International Publishing, Cham, 261–275. https://doi.org/10.1007/978-3-031-36033-6_17

[17] Peter Dalsgaard, Michael Mose Biskjaer, and Jonas Frich. 2023. Capturing and revisiting ideas in the design process: A longitudinal technology probe study. *Design Studies* 88 (Sept. 2023), 101200. https://doi.org/10.1016/j.destud.2023.101200

[18] Robert M. Emerson, Rachel I. Fretz, and Linda L. Shaw. 2011. *Writing Ethnographic Fieldnotes, Second Edition.* University of Chicago Press. Google-Books-ID: k83BlbBHubAC.

[19] Douglas C. Engelbart. 1995. Toward augmenting the human intellect and boosting our collective IQ. *Commun. ACM* 38, 8 (Aug. 1995), 30–32. https://doi.org/10.1145/208344.208352

[20] Thomas Erickson. 1996. The design and long-term use of a personal electronic notebook: a reflective analysis. In *Proceedings of the SIGCHI conference on Human factors in computing systems common ground - CHI '96*. ACM Press, Vancouver, British Columbia, Canada, 11–18. https://doi.org/10.1145/238386.238392

[21] Michele C. Everett and Margaret S. Barrett. 2012. "Guided tour": a method for deepening the relational quality in narrative research. *Qualitative Research Journal* 12, 1 (April 2012), 32–46. https://doi.org/10.1108/14439881211222714 00000.

[22] A. Faatz, B. Zimmermann, and E. Godehardt. 2009. Luhmann's Slip Box – What can we Learn from the Device for Knowledge Representation in Requirements Engineering?. In *2009 Second International Workshop on Managing Requirements Knowledge*. 1–3. https://doi.org/10.1109/MARK.2009.6

[23] Sarah Faisal, Simon Attfield, and Ann Blandford. 2009. A Classification of Sensemaking Representations. In *CHI 2009 Workshop on Sensemaking*.

[24] Julien Gori, Han L. Han, and Michel Beaudouin-Lafon. 2020. FileWeaver: Flexible File Management with Automatic Dependency Tracking. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology (UIST '20)*. Association for Computing Machinery, New York, NY, USA, 22–34. https://doi.org/10.1145/3379337.3415830

[25] Frank Halasz, G. 1988. Reflections on NoteCards: seven issues for the next generation of hypermedia systems. *Commun. ACM* 31, 7 (July 1988), 836–852. https://doi.org/10.1145/48511.48514

[26] Frank G. Halasz, Thomas P. Moran, and Randall H. Trigg. 1986. Notecards in a nutshell. *ACM SIGCHI Bulletin* 17, SI (May 1986), 45–52. https://doi.org/10.1145/30851.30859

[27] Han L. Han, Junhang Yu, Raphael Bournet, Alexandre Ciorascu, Wendy E. Mackay, and Michel Beaudouin-Lafon. 2022. Passages: Interacting with Text Across Documents. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–17. https://doi.org/10.1145/3491102.3502052

[28] Jeffrey Heer. 2019. Agency plus Automation: Designing Artificial Intelligence into Interactive Systems. *Proceedings of the National Academy of Sciences* 116, 6 (Feb. 2019), 1844–1850. https://doi.org/10.1073/pnas.1807184115

[29] Amber Horvath, Brad Myers, Andrew Macvean, and Imtiaz Rahman. 2022. Using Annotations for Sensemaking About Code. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*. Association for Computing Machinery, New York, NY, USA, 1–16. https://doi.org/10.1145/3526113.3545667

[30] Eric Horvitz. 1999. Principles of Mixed-Initiative User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. ACM, New York, NY, USA, 159–166. https://doi.org/10.1145/302979.303030

[31] Hyeonsu B Kang, Tongshuang Wu, Joseph Chee Chang, and Aniket Kittur. 2023. Synergi: A Mixed-Initiative System for Scholarly Synthesis and Sensemaking. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. Association for Computing Machinery, New York, NY, USA, 1–19. https://doi.org/10.1145/3586183.3606759

[32] Alison Kidd. 1994. The marks are on the knowledge worker. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '94)*. Association for Computing Machinery, New York, NY, USA, 186–191. https://doi.org/10.1145/191666.191740

[33] Amy. Ko, Robin Abraham, Laura Beckwith, Alan Blackwell, Margaret Burnett, Martin Erwig, Chris Scaffidi, Joseph Lawrance, Henry Lieberman, Brad Myers, Mary Beth Rosson, Gregg Rothermel, Mary Shaw, and Susan Wiedenbeck. 2011. The State of the Art in End-user Software Engineering. *ACM Comput. Surv.* 43, 3 (April 2011), 21:1–21:44. https://doi.org/10.1145/1922649.1922658

[34] Andrew Kuznetsov, Joseph Chee Chang, Nathan Hahn, Napol Rachatasumrit, Bradley Breneisen, Julina Coupland, and Aniket Kittur. 2022. Fuse: In-Situ Sensemaking Support in the Browser. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*. Association for Computing Machinery, New York, NY, USA, 1–15. https://doi.org/10.1145/3526113.3545693

[35] Michael Xieyang Liu, Andrew Kuznetsov, Yongsung Kim, Joseph Chee Chang, Aniket Kittur, and Brad A. Myers. 2022. Wigglite: Low-cost Information Collection and Triage. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*. Association for Computing Machinery, New York, NY, USA, 1–16. https://doi.org/10.1145/3526113.3545661

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

[37] Catherine C. Marshall and Frank M. Shipman. 1993. Searching for the missing link: discovering implicit structure in spatial hypertext. In *Proceedings of the fifth ACM conference on Hypertext (HYPERTEXT '93)*. Association for Computing Machinery, New York, NY, USA, 217–230. https://doi.org/10.1145/168750.168826

[38] John S Morabito and Joel Chan. 2021. Managing Context during Scholarly Knowledge Synthesis: Process Patterns and System Mechanics. In *Creativity and Cognition (C&C '21)*. Association for Computing Machinery, New York, NY, USA, 1–5. https://doi.org/10.1145/3450741.3465244 00000.

[39] T. H. Nelson. 1965. Complex information processing: a file structure for the complex, the changing and the indeterminate. In *Proceedings of the 1965 20th national conference (ACM '65)*. Association for Computing Machinery, Cleveland, Ohio, USA, 84–100. https://doi.org/10.1145/800197.806036

[40] Syavash Nobarany, Mona Haraty, and Brian Fisher. 2012. Facilitating the Reuse Process in Distributed Collaboration: A Distributed Cognition Approach. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*

(CSCW '12). ACM, New York, NY, USA, 1223–1232. https://doi.org/10.1145/2145204.2145388

[41] Gerard Oleksik, Hans-Christian Jetter, Jens Gerken, Natasa Milic-Frayling, and Rachel Jones. 2013. Towards an information architecture for flexible reuse of digital media. In *Proceedings of the 12th International Conference on Mobile and Ubiquitous Multimedia (MUM '13)*. Association for Computing Machinery, New York, NY, USA, 1–10. https://doi.org/10.1145/2541831.2541866

[42] Gerard Oleksik, Natasa Milic-Frayling, and Rachel Jones. 2012. Beyond data sharing: artifact ecology of a collaborative nanophotonics research centre. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*. ACM, Seattle Washington USA, 1165–1174. https://doi.org/10.1145/2145204.2145376

[43] Srishti Palani, Zijian Ding, Austin Nguyen, Andrew Chuang, Stephen MacNeil, and Steven P. Dow. 2021. CoNotate: Suggesting Queries Based on Notes Promotes Knowledge Discovery. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3411764.3445618

[44] Peter Pirolli and Stuart Card. 2005. The Sensemaking Process and Leverage Points for Analyst Technology as Identified through Cognitive Task Analysis. In *Proceedings of International Conference on Intelligence Analysis*, Vol. 5. 2–4.

[45] Yvette Pyne and Stuart Stewart. 2022. Meta-work: how we research is as important as what we research. *British Journal of General Practice* 72, 716 (March 2022), 130–131. https://doi.org/10.3399/bjgp22X718757 Publisher: British Journal of General Practice Section: Analysis.

[46] Xin Qian, Katrina Fenlon, Wayne G. Lutters, and Joel Chan. 2020. Opening Up the Black Box of Scholarly Synthesis: Intermediate Products, Processes, and Tools.. In *Proceedings of ASIST 2020*.

[47] Simon Rowberry. 2023. Historiographies of Hypertext. In *Proceedings of the 34th ACM Conference on Hypertext and Social Media*. ACM, Rome Italy, 1–10. https://doi.org/10.1145/3603163.3609038

[48] Daniel M. Russell, Mark J. Stefik, Peter Pirolli, and Stuart K. Card. 1993. The Cost Structure of Sensemaking. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems (CHI '93)*. ACM, New York, NY, USA, 269–276. https://doi.org/10.1145/169059.169209

[49] Saqib Saeed, Volkmar Pipek, Markus Rohde, Christian Reuter, Aparecido Fabiano Pinatti De Carvalho, and Volker Wulf. 2019. Nomadic Knowledge Sharing Practices and Challenges: Findings From a Long-Term Case Study. *IEEE Access* 7 (2019), 63564–63577. https://doi.org/10.1109/ACCESS.2019.2916903

[50] Raymond Scupin. 1997. The KJ Method: A Technique for Analyzing Data Derived from Japanese Ethnology. 56, 2 (1997), 233–237. https://dx.doi.org/10.17730/humo.56.2.x335923511444655

[51] Moushumi Sharmin, Brian P. Bailey, Cole Coats, and Kevin Hamilton. 2009. Understanding knowledge management practices for early design activity and its implications for reuse. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Boston MA USA, 2367–2376. https://doi.org/10.1145/1518701.1519064 00047.

[52] Ben Shneiderman. 1989. Reflections on authoring, editing, and managing hypertext. In *The society of text: Hypertext, hypermedia, and the social construction of information*, Edward Barrett (Ed.). MIT Press, Cambridge, MA, US, 115–131. Publisher: MIT Press Cambridge, MA.

[53] Kenneth Strike and George Posner. 1983. Types of synthesis and their criteria. In *Knowledge Structure and Use*, S Ward and L Reed (Eds.). Temple University Press, Philadelphia, PA.

[54] Sangho Suh, Bryan Min, Srishti Palani, and Haijun Xia. 2023. Sensecape: Enabling Multilevel Exploration and Sensemaking with Large Language Models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. ACM, San Francisco CA USA, 1–18. https://doi.org/10.1145/3586183.3606756

[55] Aurélien Tabard, Wendy E. Mackay, and Evelyn Eastmond. 2008. From individual to collaborative: the evolution of prism, a hybrid laboratory notebook. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work (CSCW '08)*. Association for Computing Machinery, New York, NY, USA, 569–578. https://doi.org/10.1145/1460563.1460653

[56] 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 (CHI '04)*. ACM, New York, NY, USA, 415–422. https://doi.org/10.1145/985692.985745

[57] Jaime Teevan, William Jones, and Robert Capra. 2008. Personal Information Management (PIM) 2008. *SIGIR Forum* 42, 2 (Nov. 2008), 96–103. https://doi.org/10.1145/1480506.1480524

[58] Leslie Thomson. 2015. The guided tour technique in information science: Explained and illustrated. *Proceedings of the Association for Information Science and Technology* 52, 1 (2015), 1–5. https://doi.org/10.1002/pra2.2015.1450520100135

[59] Randall H. Trigg and Peggy M. Irish. 1987. Hypertext habitats: experiences of writers in NoteCards. In *Proceeding of the ACM conference on Hypertext - HYPERTEXT '87*. ACM Press, Chapel Hill, North Carolina, United States, 89–108. https://doi.org/10.1145/317426.317435

[60] Gary Wolf. 1995. The Curse of Xanadu. *Wired* (June 1995). https://www.wired.com/1995/06/xanadu/

[61] Joost F Wolfswinkel, Elfi Furtmueller, and Celeste P M Wilderom. 2013. Using grounded theory as a method for rigorously reviewing literature. *European Journal of Information Systems* 22, 1 (Jan. 2013), 45–55. https://doi.org/10.1057/ejis.2011.51 Publisher: Taylor & Francis.

[62] Zhang Pengyi and Soergel Dagobert. 2014. Towards a Comprehensive Model of the Cognitive Process and Mechanisms of Individual Sensemaking. *Journal of the Association for Information Science and Technology* 65, 9 (Aug. 2014), 1733–1756. https://doi.org/10.1002/asi.23125