Exploring Team-Sourced Hyperlinks to Address Navigation Challenges for Low-Vision Readers of Scientific Papers

SOYA PARK, Massachusetts Institute of Technology, USA
JONATHAN BRAGG, Allen Institute for Artificial Intelligence, USA
MICHAEL CHANG, Stanford University, USA
KEVIN LARSON, Microsoft Research, USA
DANIELLE BRAGG, Microsoft Research, USA

Reading academic papers is a fundamental part of higher education and research, but navigating these information-dense texts can be challenging. In particular, low-vision readers using magnification encounter additional barriers to quickly skimming and visually locating information. In this work, we explored the design of interfaces to enable readers to: 1) navigate papers more easily, and 2) input the required navigation hooks that AI cannot currently automate. To explore this design space, we ran two exploratory studies. The first focused on current practices of low-vision paper readers, the challenges they encounter, and the interfaces they desire. During this study, low-vision participants were interviewed, and tried out four new paper navigation prototypes. Results from this study grounded the design of our end-to-end system prototype Ocean, which provides an accessible front-end for low-vision readers, and enables all readers to contribute to the backend by leaving traces of their reading paths for others to leverage. Our second study used this exploratory interface in a field study with groups of low-vision and sighted readers to probe the user experience of reading and creating traces. Our findings suggest that it may be possible for readers of all abilities to organically leave traces in papers, and that these traces can be used to facilitate navigation tasks, in particular for low-vision readers. Based on our findings, we present design considerations for creating future paper-reading tools that improve access, and organically source the required data from readers.

CCS Concepts: • **Human-centered computing → Accessibility design and evaluation methods**; *Accessibility systems and tools.*

Additional Key Words and Phrases: accessibility, low-vision users, accessible scholarly, tools for researchers, reading interface

ACM Reference Format:
Soya Park, Jonathan Bragg, Michael Chang, Kevin Larson, and Danielle Bragg. 2022. Exploring Team-Sourced Hyperlinks to Address Navigation Challenges for Low-Vision Readers of Scientific Papers. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 516 (November 2022), 23 pages. https://doi.org/10.1145/3555629

1 INTRODUCTION

Navigating academic papers is a fundamental part of the research process and higher education. Research findings are typically published in PDF documents, which include a body of text combined with tables, figures, and supplementary materials [37]. Other researchers and students read these PDFs to gain background knowledge, learn about and build upon related work, and find answers to

Authors' addresses: Soya Park, soya@mit.edu, Massachusetts Institute of Technology, USA; Jonathan Bragg, Allen Institute for Artificial Intelligence, USA, jbragg@allenai.org; Michael Chang, Stanford University, USA; Kevin Larson, Microsoft Research, USA; Danielle Bragg, Microsoft Research, USA.

2573-0142/2022/11-ART516
https://doi.org/10.1145/3555629

specific questions. In particular, academic readers may navigate between different parts of a paper to check citations, to re-read parts that they did not understand, and find particular information (e.g., study methodology, terminology definitions, "as mentioned above" contents).

Navigating academic papers can be a challenge. Locating desired textual content can be difficult, as academic content is typically dense, and writing style and paper structure vary. Finding information in a referenced figure or table can also be difficult, especially when figures and tables are small, illegible, or poorly labelled. Text sometimes provides links to citations or specific paper sections, figures, or tables, but these links are typically one-directional, which makes returning to the body of the text difficult. Supporting tools for automatically extracting desired information are limited, as the PDF format is difficult to automatically parse. Once desired material is found, there is also limited support for bookmarking the content, preserving links identified between sections, and sharing these navigation hooks with colleagues and other readers. This results in a large amount of duplicate work across readers.

Many navigation difficulties are amplified for readers with low vision. Visual impairments are pervasive – worldwide, there are more than 2.2 billion people with a visual impairment [25], and most have some vision (i.e., are not blind). Readers with low-vision often prefer to read visually, rather than with a screen reader or text-to-speech software. Low-vision conditions are diverse, so to supplement their vision, low-vision users juggle a suite of assistive technologies [35]. To do this, they typically use high magnification, which reduces the viewport – the amount of visible content. As a result, they need to constantly adjust their viewport either by moving their cursor or scrolling. This constant zooming and panning makes navigation more difficult – from gaining a high-level view of a figure or paper structure, to skimming and locating desired contents. Small screens (e.g., smartphones) similarly reduce the viewport and exacerbate navigation challenges for sighted readers as well.

Previous research suggests navigation links as a promising method for navigation within a complex document [23]. Despite recognizing the benefit of the links, it is still unclear how to *generate* such links or automate creation of a wider variety of links. Within the main text, many PDFs provide links to references and paper elements including section titles, figures, and tables. These links are typically input manually by the authors (e.g., using LaTeX). Recently, researchers have tackled the problem of automatically linking scientific terms with definitions (e.g., [14]). However, there is a broader range of contents that readers jump between, which AI is not currently able to link. For example, readers may jump back and forth between a results section and relevant parts of an experiment design; between a figure and the text describing it; or between a discussion point and a referenced result.

In this work, we explore academic-paper reading interfaces that facilitate navigation, in particular for low-vision readers. We engaged in a sequence of two studies to explore this space. We began with a first study focused on understanding how low-vision readers navigate academic papers, what problems they encounter, and what types of paper navigation features they might desire. The user feedback from this study enabled us to then build a complete navigation system with an interface that aligns with low-vision reader preferences. To enable desired navigation features, the system leverages data from other readers. Our second study then focused on understanding the experience of groups of low-vision and sighted users contributing to and using such a tool. We conclude our work by consolidating design guidelines for creating effective paper navigation interfaces, in particular for low-vision users, and for leveraging navigation interactions to improve navigation for future readers.

In our first study, to ground our system design, we interviewed six low-vision readers who read academic papers daily, and asked them to use a set of design probe interfaces that augment navigation within an academic paper, based on hard-coded data. We identified their current strategies of navigation during reading academic papers, and asked them how and whether each interface helps their navigation. We found that low-vision readers may desire features that they can initiate throughout the paper, do not alter the original narrative of a paper, and allow them to read through the paper and unpack the contents afterward.

In our second study, to better explore potential system usage, we conducted a two-week-long lab study using a design probe system that not only provides user-facing navigation elements, but also bootstraps the required data from other users. Our prototype probe, which we call Ocean, was designed based on accessibility requirements from the first study, and enables users to leave anchors and links that they themselves use for navigating to different parts of a paper. By sharing these anchors and links across users, the system organically crowdsources navigation links that AI cannot currently automate. Participants were invited in teams and asked to create and use shared links for various navigation tasks. Through the study, we found that both low-vision and sighted users took advantage of their own or team members' links for the tasks and generally felt that they were able to produce links with low effort.

In summary, the contributions of the paper are as follows:

- Need-finding interviews to better understand low-vision users' needs for academic paper reading and navigation and the barriers they encounter therein. We learned that many low-vision users may need better interfaces to navigate information in academic papers. In particular, low-vision users may benefit from an interface which preserves their viewport and does not rely on imperfect automated accessibility tools.
- An exploratory interface for accessible academic paper reading, Ocean, designed to meet the needs of low-vision users. To help provide links that automated systems may be unable to generate, the interface allows users to leave navigation traces as they read and navigate papers. Future readers can then reuse and benefit from these navigation traces or links.
- A case study using Ocean. Results suggest that traces can be left by both low-vision and sighted users, and that navigation traces created using Ocean may enable both low-vision and sighted readers to expedite certain navigation and information-seeking tasks.
- Design considerations for creating future paper navigation tools that improve access for low-vision readers. These considerations build on our findings and on prior work, and include insights about accessible front-end design, and the potential for incorporating data generated by users engaged in tasks that benefit their own reading experience.

2 RELATED WORK

We sought a new design of reading interface navigation methods. We leveraged navigation traces left by previous readers, so future readers can use them for navigating information. In this way, all of the participants (both low vision and sighted) are willing to contribute, not primarily because of altruistic reasons but to help their reading, which eventually improve the accessibility of the complex document for low-vision readers. Here, we reviewed previous work on interfaces for readers with visual impairments and academic-paper reading interfaces. Lastly, we shared work about improving the accessibility of a workspace with the help and volunteering of peers.

2.1 Reading interfaces for people with visual impairments

Researchers have explored the challenge of reading on screens with visual impairments and interfaces that targeted their needs. Blind users typically rely on screen readers, which read aloud text and interface elements, when they read. These tools support reading full texts or sections sequentially, but do not easily support navigation or skimming [1]. They rely on headers, tags, and other elements to understand text and interface structures, and to enable users to navigate content.

For example, to consume images non-visually, alt-text, which is a description of an image, may be embedded in a PDF. These descriptions are created through interpretation by sighted people or conversion of contents into a simplified form [34]. However, alt-text is often missing, and typically lacks details and structure needed to understand images without viewing them [22]. Moreover, screen readers provide a separate and often worse user experience for visually impaired users, and do not address the need for a shared and equitable experience between readers who are sighted and visually impaired [29].

Unlike blind users, low-vision users have limited sight, which may enable them to read visually. Many low-vision readers use screen magnification or color inversion (e.g. prefering white text on black background), with or without screen readers [19, 35]. Nonetheless, they face a unique set of challenges when reading online. Wu et al. highlight the difficulties of reading complex documents with low vision: 1) detailed and small fonts can decrease legibility, especially when the interface does not wrap the text, 2) checking and following references across articles is challenging with text-to-speech or large magnification, and 3) due to varied text layouts, technology cannot accurately comprehend or parse the article in a PDF format [38].

Despite such challenges of reading with a visual impairment, there are social barriers to using assistive technology. People with visual impairments often do not want to disclose their disability, and using large magnification or a screen reader can be conspicuous [35]. Tools used by organizations are often determined by people without visual impairments, and accessibility is often not considered [7]. Inaccessible interfaces lead visually impaired users to juggle multiple tools and create *segregated* experiences between sighted and visually impaired users [20, 29]. In this study, we aimed to explore interfaces for academic paper navigation with a focus on an accessible, unified experience for sighted and low-vision users.

2.2 Academic reading tools

Academic reading tools for researchers have been broadly developed and studied, some of which low-vision researchers can benefit from. Some work has focused on facilitating discussion of papers by readers, for example by embedding discussion threads within academic papers. The result indicated that such situated discussion can lead to a better understanding of the contents and facilitate learning between peers [39, 40]. Other work has focused on enabling researchers to find relevant content, for example by tagging papers with richer attributes such as topics and semantic classes [9, 10, 21, 24, 30]. With these enriched attributes, researchers were able to search, filter or visualize numerous articles.

Most relevant to our current work is prior work focused on local navigation within a paper. In particular, previous work has created links between visualizations and referencing texts in academic articles [18] and between definitions of terms and their usage in a paper [14]. Each of these projects used AI to detect and add a very specific type of link. However, this type of approach is not currently viable for a variable and poorly defined range of links or navigation jumps that users engage in (e.g., jumping between results and methodology, or jumping from a mention of a tool in the abstract directly to the description of the tool in the main text). In this work, we explored a navigation interface that enables the crowd of users to leave navigation traces, which can be used for their or others' information seeking.

While rich sets of tools and approaches have been explored, limited progress has been made on the accessibility of scientific literature; Wang et al. [37] found that academic papers were highly inaccessible for researchers with visual impairments and investigated the challenges and practices of visually-impaired researchers and proposed a system that converts academic papers in an accessible format for the researchers, but without the general navigation aides we study in this work.

In our work, we aimed to augment previous work and explored the design space of information navigation in academic papers for low-vision researchers. We first conducted a need-finding interview study accompanied by four prototypes and a field study of a reading interface. We then concluded our work by consolidating design guidelines from both studies for reading interfaces for low-vision researchers.

2.3 Collaborative accessibility tools

Leveraging friends' help (friendsourcing) has been proposed as a reliable and scalable method to improve accessibility of online contents, as well as overcome shortcomings (i.e. the issue of quality control) of AI techniques and crowdworkers [28]. Unlike employing AI or crowdsourcing, friendsourcing accessibility is often altruistically motivated (i.e. helping a blind friend) [3] and can improve mutual bonds of caregivers and people with disabilities [5, 32]. Nevertheless, many people with visual impairments and other disabilities perceive this type of help as a burden and may feel indebted to collaborators or helpers [27]. Previous work reveals that low-vision people in particular tend to be independent and reluctant to ask for help from others [35]. This trait suggests a new opportunity in the design space for collaboratively-seeded tools for low-vision users, focused on how to elicit inputs or help from others without this social cost.

Organic crowdsourcing [17] is a technique that has been applied in other domains, to collect useful data from people while they do tasks that they normally do. Like in traditional crowdsourcing tasks, people accomplish small units of work that contribute to a larger effort. However, in organic crowdsourcing, contributors do not inconvenience themselves or take on tasks that they are not intrinsically motivated to pursue. The result is improved scalability, and reduction or elimination of debt to contributors – whether monetary or social. This framework has been applied to advance a host of scientific discoveries [15, 31], for social experiments [26] and for some accessibility applications (e.g. collecting sign language data [4]). Our work can be viewed as leveraging organic crowdsourcing to seed paper navigation tools, in particular in support of low-vision readers, scalably and without burdening contributors.

3 STUDY 1: LOW-VISION DESIGN PREFERENCES FOR NAVIGATION INTERFACES

We began by investigating navigation needs and interface preferences of low-vision readers, with the goal of collecting design insights for our final interface. We conducted this need-finding study first, to enable us to subsequently design and evaluate a complete system that aligns with low-vision reader preferences (see Study 2). In this first study, we conducted semi-structured interviews with individuals with low vision to better understand: 1) low-vision readers' current practices for reading academic papers, 2) their navigation challenges, and 3) the usefulness and usability of different interfaces for navigation.

3.1 Participants

We recruited participants via posting on mailing lists of an IT company and a private university. Each interviewee was compensated $65 for their time. Participants (2 females, 4 males, age=22-53) were researchers, students, project managers, or admins who self-identified as having low vision and who visually read academic papers regularly. Two of the participants combined text-to-speech with visual reading.

3.2 Protocol

Each participant was interviewed by the first author over Zoom, which took about one hour in total. We first asked participants to share their usual way of reading academic papers and the challenges they encounter. Then, we asked them to share their screen with us, read an academic paper using the

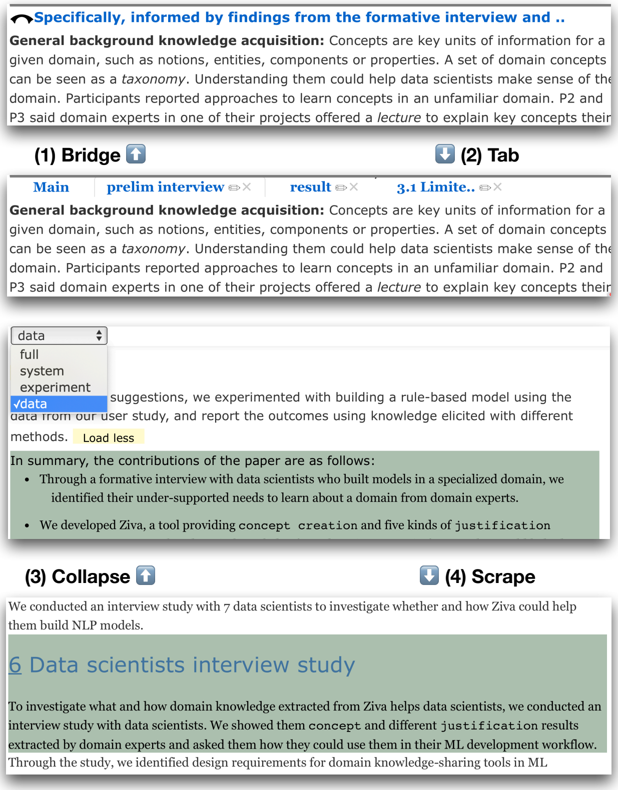

Specifically, informed by findings from the formative interview and ..

General background knowledge acquisition: Concepts are key units of information for a given domain, such as notions, entities, components or properties. A set of domain concepts can be seen as a *taxonomy*. Understanding them could help data scientists make sense of the domain. Participants reported approaches to learn concepts in an unfamiliar domain. P2 and P3 said domain experts in one of their projects offered a *lecture* to explain key concepts their

General background knowledge acquisition: Concepts are key units of information for a given domain, such as notions, entities, components or properties. A set of domain concepts can be seen as a *taxonomy*. Understanding them could help data scientists make sense of the domain. Participants reported approaches to learn concepts in an unfamiliar domain. P2 and P3 said domain experts in one of their projects offered a *lecture* to explain key concepts their

suggestions, we experimented with building a rule-based model using the data from our user study, and report the outcomes using knowledge elicited with different methods.

In summary, the contributions of the paper are as follows:

- Through a formative interview with data scientists who built models in a specialized domain, we identified their under-supported needs to learn about a domain from domain experts.
- We developed Ziva, a tool providing `concept creation` and five kinds of `justification`

We conducted an interview study with 7 data scientists to investigate whether and how Ziva could help them build NLP models.

6 Data scientists interview study

To investigate what and how domain knowledge extracted from Ziva helps data scientists, we conducted an interview study with data scientists. We showed them `concept` and different `justification` results extracted by domain experts and asked them how they could use them in their ML development workflow. Through the study, we identified design requirements for domain knowledge-sharing tools in ML

Fig. 1. The four web prototypes presented in the needfinding interviews (Study 1). (1) Bridge: users can navigate through hyperlinks connecting between two different parts of a paper (bridge icon in top left is a clickable anchor). (2) Tab: users can create multiple tabs within a paper (first tab is selected). (3) Collapse: users can filter texts by the class of information. (4) Scrape: users can re-arrange information in a paper and load relevant parts of the paper right below of where they were reading

ACM web reader and use their normal approach of reading a paper. Finally, we asked participants to use our prototypes, and discussed their experience, where the order of presented interface is shuffled for each interviewee. We sought to explore design spaces and use our prototypes as design probes [36]. During their reading, we asked them to think aloud during each step, for example to describe the part of the paper they were thinking of reading and how and why they would accomplish the navigation. The first author went through the interview transcripts and coded them for themes using an open-coding protocol. Through multiple iterations along with periodic discussions with the rest of the research team, the coding led to 34 codes, from which the following major themes were selected. Because of the low number of interviewees, our interview findings should be regarded as indicative.

3.3 Design Probe Prototypes

Interviewees were introduced to the following four interfaces (Bridge, Tab, Scrape, and Collapse; see Figure 1) in a randomized order. We developed these prototypes through a participatory design process with a low-vision researcher. We sought to create a variety of interfaces that help navigate information within an academic paper and enable easy access to relevant information, without designing for how the required data would be created (e.g., links between relevant parts). For the user study, we hard-coded the required data for four papers at the intersection of human-computer interaction and machine learning, which was within the research expertise of all participants. We study how to scale the creation of this required data in our subsequent study. The four prototypes were built in built on the HTML format, as the format is more accessible to readers with visual impairments [37].

The Bridge interface is inspired by the literature on hypertexts, which enable improved navigation with non-linear access [23]. This interface incorporates hyperlinks in a paper. Different from many prior systems, this interface provides bi-directional links that let a user jump back and forth between two parts of the paper. The links can be anchored not only at headings but also within a sentence (e.g., where a notable concept is defined). When a user clicks either of the link's two anchors, the interface jumps to the other end of the link. The interface displays the snippet of the other end of the link at the top toolbar to preview the other end before users jump. For example, if there is a citation that a user wants to investigate, they can click the link to jump to the corresponding reference, and later return to the main text they were reading by clicking the snippet of another end (or pressing the keyboard shortcut “b”).

Inspired by tabbed browsing, the Tab interface lets a user open a new tab whenever they want to view different parts of the paper. The new tab displays the same paper, but is scrolled to a different part. Like the Bridge interface, a tab can be created from pre-populated anchors throughout the paper. The name of the new tab is initialized with the section name that the user is initially at, but they can also rename tabs. They can move between tabs by using shortcut numbers – 1 is the first tab, 2 is the second tab, and so on. Users can also duplicate a tab by right-clicking the tab.

Academic papers include various classes (e.g., system description, study methods) of information [24]. The Collapse interface makes excerpts of related text available in expandable sections interleaved with the main text. If the user wants to continue reading beyond the excerpt, they can click the “load more” button. They can also load less by clicking the “load less” button or right-clicking the text to access the context menu.

The Scrape interface re-arranges the order of paragraphs and sections in a paper as a user needs, so the user can read information that they are interested in first, regardless of composition of the paper. If they want to read more about a certain part of the paper, they can bring the linked part of the paper into the current text flow, on demand. By clicking a “load more” button near an anchor, the interface attaches the relevant part of the paper below the button. If they want to close the expanded text, they can right-click the text and close it.

3.4 Results

All interviews were transcribed, and analyzed by axial coding of the 6 interview transcripts following an open-coding protocol. The following subsections outline the findings of the study, including participants' current set-up and practices for reading academic papers and their desired functionality for their reader interface.

3.4.1 Current practices for reading academic papers. Interviewees employ different set-ups to read academic papers. They used browser magnification (2 out of 6), browser magnification in a color-inverted large monitor while reading in close proximity to the screen (1), browser magnification with text-to-speech (2), and a camera of their mobile device as a magnifier (1) for reading academic papers. None of the participants used system magnification since it “breaks the format” of the paper and the text does not wrap within their viewport.

3.4.2 Strategies for and Challenges of Reading Academic Papers with Low Vision.

Academic papers are information-packed and can be difficult to read with a limited viewport. Interviewees said academic papers are often difficult to understand because the information load is overwhelming. They pointed out the length of papers as one of the reasons: "academic papers… really tend to put lot of information in [many] pages, so that's one thing which I find [makes] the reading typically so [challenging]." They added that navigating long texts especially impacts readers with low vision because their viewport – the amount of text that fits on screen – is limited due to magnification. As a result, reading text and parsing figures and tables requires panning around, which consumes a substantial amount of time and energy: "Because the diagrams are big [and] large… [they] go through a full page so it's difficult to… get one single view of the whole diagram." This also happens for the texts in viewers that do not wrap the text (e.g., for PDF). This constant panning decreases their reading speed and impacts understanding of the paper: "But my rate of reading was incredibly slow [..] And [..] because I was reading so slowly as well, I would lose track of the central idea, so it was just not working for me."

Readers with low vision often encounter difficulties of paper navigation and find it difficult to orient themselves while reading academic papers. Some interviewees jump to related sections if they are unable to unpack the part that they are reading. For example, as introduction sections often involve a substantial amount of information, readers needed to jump to the relevant section before they finish the introduction. "I need to stop [reading] in the first page and then go to the next. If I can't get a full view in academic paper [then I need to] read the system [section]." However, jumping to other sections comes with a great cost for low-vision readers. Since they have a small viewport, it is difficult to locate the part they want to go to and come back to the part they were reading. For example, there is no easy way to quickly check out a citation and come back to the current text: "So I don't have an easy way to get back to where this was." A few interviewees (3/6) remarked that they essentially "gave up" on navigating within a paper because they found finding the relevant part and coming back afterwards too challenging. Rather, they kept reading and hoped that the contents would become clear soon: "Not understanding everything 100% in the abstract but maybe if it's somewhere like methodology or something, you know, a few sections in I would probably stick around until I understand it. But I would like to know [..] what are the results of this paper like how this really [..] works."

Interviewees developed varied strategies to help them navigate papers. First, they searched the document by a keyword that they would like to learn more about: "What I'd probably do is do a `Control-home` to take myself to the top and then I'm going to probably do `Control+F` [..] then enter and now it's taking me to [citation] 83." In order to be able to return to their current reading location, they memorized a reasonably unique nearby word, which they could later search for. Interviewees said that search is especially useful for them since search is universally accessible by keyboard: "Visually disabled people like me, mainly use keyboards to use the computer." However, this strategy is imperfect since keyword match is brittle, for example when searching for paraphrased terms or terms that are hyphenated and split between lines. Search may also require sifting through many matches, and increase cognitive load.

Another strategy participants employed is to utilize section titles and document structures, echoing previous work that showed visually impaired readers rely heavily on headings [37]. Headings often have bigger fonts, which are more visually distinguishable and easier to read, and may be hyperlinked in the original document from a table of contents. Since the structure of academic papers is fairly uniform, at least within an academic discipline, low vision readers can guess which section likely contains the desired information, jump to that section, and start reading. However, finding the desired content within a section still requires reading a substantial portion of the section, and readers may incorrectly guess which section contains the desired information.

Furthermore, scrolling to a desired section is still cumbersome and time-consuming unless the viewer has a table of contents with links. It can also be difficult to orient oneself within a paper, and determine which section they are currently reading. An interviewee said that some PDF viewers show the current section name if they hover over the scroll bar, which helps to orient the reader: "I always prefer a kind of a magnifying thing like when I hover over this part [..] the sidebar which basically tells like which section it is."

Lastly, they highlighted important sentences or created notes on the PDF version of papers to anchor parts of papers that they found meaningful. This helps them not only understand the paper but also create custom anchors for returning later. An interviewee said: "So once I read the paper, [I'm] going through like [and] I just make some points on my notes, so I would make out like [..] what are the key concepts of this project, especially like this. I'd like to note down this or that; I will go to the introduction part and read the introduction." Note-taking may present additional difficulties to low-vision readers who incorporate text-to-speech or screen readers in their reading. They said when listening to text, they are reluctant to pause the audio and mark the part since it interrupts their reading. An interviewee who uses text-to-speech remarked that they ended up not annotating papers and "pay the price for barely doing it" after finishing reading the paper: "I just said it's not worth the time and I just keep going. But then obviously I end up with non-highlighted PDFs and no notes, which is a horrible way of reading a paper. Basically... fully understand that people want reading and if it's not working, that is challenging."

General skepticism towards assistive technology. Since low-vision researchers are typically still able to visually read, many forgo assistive technology for academic paper reading. All of the interviewees remarked that they had explored tools for reading academic papers, such as document converters for converting to or from PDF, web viewers, and text-to-speech, and that the assistive tools "screwed up significantly." This made them anticipate malfunctions when using these tools, which resulted in them having to constantly supervise system performance, which is a large distraction. In the worst case, they abandoned the tool. An interviewee said: "The issue with these sorts of [assistive] software is that they're not like reliable all the time, like across papers. This biases against using the software, because you have this representation in your head that "Okay, the software isn't [going to be] that great because it hasn't been great in the past [..] I feel like it would be really frustrating [..] if it doesn't work for like two of the three things that I click, [..] I probably would stop using the tool immediately from there."

3.4.3 *Types of navigation within an academic paper.*

During our need-finding interviews, participants shared various types of navigation that they perform when reading an academic paper. We consolidated these types of navigation into four sublists (Tables 1–4), based on the following categories of motivations. Researchers sought additional information relating to the part they were reading due to interest or curiosity (Table 1), difficulty in comprehension (Table 2), desire to compare two analogous concepts (Table 3), or failure to recall the meaning of a concept (Table 4).

3.4.4 *Preference for Different Navigation Methods.*

After gaining understanding of participants' challenges while reading academic papers, we then presented them with our prototypes and asked for their thoughts and preferences. Participants had varied reasons for their preferences.

Preserving viewports and accessing interfaces with less scrolling. One of the popular reasons behind interviewees' preferred interface is how well the interface preserves viewports. In addition, participants preferred interfaces that supported navigation without needing to pan or scroll, such as by keyboard or sticky headers. As mentioned earlier, finding and adjusting viewports happens frequently and is a cumbersome task. Hence, it is critical that interfaces do not change their viewport unexpectedly. An interviewee said that viewport preservation is important to them in general,

Table 1. Navigation types: Looking for additional information due to interest or curiosity

Description
Link between text referencing material in another core section or subsection (e.g., figure [18], table, footnote, citation)
Link to details about the upcoming concepts (backward referencing) e.g., "(further discussed in a later section)"
Link from abstract to evaluation to see how they actually measure and lead to that conclusion
Link between motivation and previous work
Link between high level system description and eval results
Link from citation and how they are used in a paper
Parts that I would like to share/discuss with others [39]
Paper discourse classes e.g., motivation, study methods, etc

Table 2. Navigation types: Looking for additional information due to difficulty in comprehension

Description
Link between poorly explained parts and explanations
Link between motivation and previous work
Link between motivation and example scenario
Go back to section intro [37]

Table 3. Navigation types: Looking for additional information due to desire to compare analogous parts/concept

Description
Link between analogous parts of the paper (e.g. two figures, two study designs, two datasets, etc.)
Link from citation and how they are used in a paper

Table 4. Navigation types: Looking for additional information due to failure to recall the meaning of a concept

Description
Link between term and definition (forward-referencing) [14]
Link between metric and metric definition [14]
Go back to section intro [37]

beyond paper reading: "The reason that [Safari is] my preferred browser is basically because you have the pinch to zoom feature like on tablets [..] In other browsers you can just make the text larger, however, that breaks the formatting of the page…Doing pinch to zoom means you can look at an area, more specifically, without really breaking the formatting of the webpage."

Accordingly, among our prototypes, interviewees preferred interfaces where viewports stay at the part they are reading and "minimize the scrolling" they need to do: "I was able to be taken back to where I was and continue reading the paper, so that was one thing which I always wanted." Interviewees also expressed their preference for interactions that do not change the viewport, such as using a keyboard or sticky headers that persist at the top of the interface: "The dropdown follows me throughout the page wherever I go so it's kind of static here, so I can just click on whatever data I want." The Tab interface was especially useful in this regard since the tabs persist at the top of the interface and users can revisit various parts of the paper with the same viewport.

Preserving original content. Interviewees highlighted a desire to have access to the original contents. They did not want the presented content to deviate from the author's narrative because they do not want to miss out or understand the paper incorrectly. They want equal access. Interviewees remarked that the Collapse and Scraping interfaces manipulate the paper's narrative, and so thought they would not be an ideal way of reading papers. One interviewee said: "The author [..] writing the paper, they had a certain flow in mind. Like they basically assume that the reader is going to be reading from [..] the beginning to the end in the same order." Especially for the Scraping interface, interviewees said that the intertwined contents interfered with the intended flow of the paper. They said that these interfaces could be useful for special cases such as skimming or finding specific information across different papers, but were not suited for in-depth academic paper reading.

Wanting to unpack loaded information. Interviewees expressed a desire to unpack loaded information and get on-demand support for understanding as they read papers. As mentioned before, due to their small viewport, finding information in academic papers can be challenging with low vision. They would like to find the information when they are stuck at some phrase of the paper. One interviewee said: "What I'm looking for is like an explanation of the phrase that I don't understand." All of the interviewees said that they have difficulties understanding the abstract and introduction of papers because the sections tend to be short and contain a substantial amount of information. For example, an interviewee said that they often want to understand better when they read an abstract and navigate from the abstract to the result section: "At this point, [..] I'll probably jump to the results first. I think that's what I would be most curious to see. And then I would sort of go back and then look at the introduction and stuff like that." Another interviewee shared a use case of the Collapse interface to unpack information and concentrate on the relevant information in each section. Navigation to earlier parts of the text is also sometimes desired, for example when the reader forgets a concept they read and wants to revisit it. Interviewees shared a preference for using the Tab or Bridge interface for the purpose: "[The paper] talked about this metric. I'm not sure what this is. So at this point, I would probably look it up."

4 STUDY 2: FIELD DEPLOYMENT OF ACCESSIBLE READING INTERFACE

While our first study focused on exploring how to design user-facing components of navigation tools for low-vision users, it did not address the challenge of how to seed such systems with the required data, nor explore how groups of readers – both low-vision and sighted – would consume the navigation features during actual reading. In our second study, we sought 1) to explore how readers can organically generate navigation links that other readers can use, and 2) to explore the usage of such a collaborative system by both low-vision and sighted users. Informed by Study 1 findings of low-vision accessibility preferences, we developed a design probe system, which we call Ocean, that both facilitates paper navigation and generates the required data from users. The prototype's front-end is designed based on the findings from Study 1, for example by providing navigation hooks that are always within the reader's viewport. Additionally, the system seamlessly generates the required data by enabling users to create, use, and share navigation links within papers. To explore the user experience, we invited teams of researchers to use the prototype to read papers and create links, to find answers to questions about the text, and to discuss their experience. In contrast to our previous study, this user study occurred over the course of a week, involved both low-vision and sighted readers, and involved deployment of a complete design probe system in a more realistic setting.

4.1 Participants

We recruited low-vision participants from our Study 1 pool and through word-of-mouth, and each low-vision researcher was asked to sign up for the study with two colleagues with whom they had previously collaborated on research projects. Each participant was compensated $122.50 USD for their time. We recruited 2 teams, spanning 6 participants (1 female, 5 males, age=24-41). Participants were graduate students, full-time researchers, or software engineers.

4.2 Protocol

To explore how Ocean may be used in the wild, we conducted a two-week-long field study. Participants were asked to read two papers using two different interfaces (Ocean and the standard ACM web interface) with their collaborators. We chose the ACM web interface as our control interface, as they provide numerous papers in an HTML format. The order in which they used the interfaces was randomized. The procedures were as follows:

1. Tutorial (1 hour): Each team was invited to a tutorial session a day before the start of the study, where they received a tutorial on how to use the Ocean prototype and tried it out on a sample paper. We also introduced them to the tasks they would conduct over the next two weeks. The primary task was to read two papers – one per week, each with a different interface – with additional tasks to ensure engagement.
2. Paper reading (1 week, 2x): The team spent one week reading a paper together with one of the interfaces (repeated with the other interface). To ensure that participants engaged early, by the end of the first day, they were asked to submit a summary of the paper, and to create at least 5 links if they were using Ocean. Creating links early helped ensure that the team had links to engage with throughout the week. At the end of the week, participants took an open-book quiz about the paper during a video call with a study facilitator and were asked to fill out a NASA-TLX survey [13] on their experience with that week's interface.
3. Exit interview (1 hour): At the end of the study, participants individually engaged in a 1-hour exit interview focused on understanding their experience and eliciting feedback on our design probe. The interview was semi-structured and guided by a list of questions attached in Appendix. Similar to Study 1, the first author conducted axial coding of the interview transcripts with an open-coding protocol. The authors then discussed and refined the codes.

The two papers that participants read during the study were chosen to be similar to one another, and chosen among recent HCI topics that are accessible to all researchers. Specifically, participants read Daskalova et al. [8] and Fujita et al. [11], which were presented in CHI 2021. These two papers are similar on several important dimensions: – they were on a similar topic, presented in the same conference session, and had a similar page length. Each team used different interfaces for each paper.

The quiz questions (provided in the Appendix) were designed to involve navigation activities in each paper, so that we could observe user engagement with our prototype's navigation features and compare how participants accomplished similar tasks in the control interface. This type of quiz task is prevalent among the literature in studies of novel systems for reading [14, 40]. The goal of the quiz is to engage users' participation, not to generate data for inferential comparisons. For each paper, four multiple choices questions were asked, each designed to require a particular navigation task (Tables 1–4). In the questions, terminologies were paraphrased and referenced figure numbers were not provided, in order to simulate imperfect recall situations where it is not feasible to find answers using `Ctrl+F`.

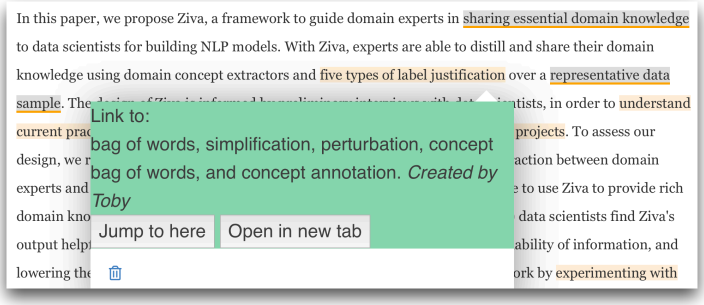

In this paper, we propose Ziva, a framework to guide domain experts in sharing essential domain knowledge
to data scientists for building NLP models. With Ziva, experts are able to distill and share their domain
knowledge using domain concept extractors and five types of label justification over a representative data
sample. The ientists, in order to understand
current prac projects. To assess our
design, we r action between domain
experts and to use Ziva to provide rich
domain kno data scientists find Ziva's
output helpf ability of information, and
lowering the ork by experimenting with

Link to:
bag of words, simplification, perturbation, concept bag of words, and concept annotation. *Created by Toby*

Fig. 2. The Ocean interface: Users can collaboratively create links and bookmarks to use them for navigation in the paper. Links and bookmarks were displayed in a different color (Links: highlighted in orange, Bookmarks: in light-grey)

4.3 Design Probe *Ocean*

Ocean was designed both to provide an accessible front-end, and to enable readers to seamlessly seed the system with navigation features between relevant parts of the paper. To ensure that our navigation interface was accessible, we largely based the front-end on our findings from Study 1. In particular, we incorporated participant preferences for interfaces that preserve viewport and keep the original contents: Because Study 1 participants valued Bridge links that enabled them to quickly spot-check information, we adopted links in Ocean. Furthermore, the participants preferred tabs since they allow users to revisit different parts of a paper, hence we adopted tabs design in Ocean as well.

To enable readers to organically seed the system with data, Ocean provides a bookmark and link authoring interface for creating an augmented reading experience without changing the flow of the original text (Figure 2). Users can highlight content to create bookmarks in any passage in a paper, which is one of the natural interactions during reading. These bookmarks can later be used to create a link which connects two different bookmarks. Using links, readers can check related information or use it to jump to another part of the paper. Bookmarks are private, i.e. only visible to their author, but links are public so other readers can use them, too. Such different privacy allows users to disclose their complete traces (i.e. links) as opposed to work-in-progress activities (i.e. bookmarks). Readers are incentivized to create links because a new link enables the creator to jump back to where they were previously reading, thus providing a direct and immediate benefit to the creator. If a user creates a link on content that appears multiple times (e.g., the name of the system) throughout a paper, Ocean automatically creates links from other occurrences of the word to the same endpoint as well.

Bookmarks and links can be used to create a tab within a paper. When a user clicks a bookmark or link, there is an option to create a new tab. The list of tabs can be accessed by right-clicking anywhere in the interface and used to transit to other tabs that the user created. In addition, we seeded the system with pre-populated links to paper elements that are currently often linked in PDFs, namely figures, tables, and citations, hence they can jump from figures, tables and citations to referenced texts, and vice versa.

To arrive at our final design, we engaged in an iterative design process with mixed-ability pilot users – a team of researchers used the system for a week and shared feedback on the interface. Our design team included both sighted and low-vision researchers with first-hand experience

Table 5. Generated types of links by participants over the deployment in each team

Types of links	**Team 1**	**Team 2**
Link between text referencing material in another core section or subsection, and that core section or subsection	58.8%	29.4%
Details about the upcoming concepts	5.9%	41.2%
Link between term and definition	-	17.6%
Link between analogous parts of the paper (e.g. two figures, two study designs, two datasets, etc.	-	11.8%
Link between poorly explained parts and explanations	29.4%	-
Link between motivation and previous work	5.9%	-

Table 6. NASA TLX responses (1- Very low, 7- Very high): Average (standard deviation)

	Mentally demanding	Successful in accomplishing tasks	Amount of work to accomplish
ACM web interface	4.17 (1.33)	3.17 (1.33)	3.83 (0.75)
Ocean	3.67 (2.16)	3.83 (1.72)	3.33 (2.06)

navigating academic papers, and experts on reading science, human-computer interaction, and human-in-the-loop computing.

We implemented our design probe as a web application built on the ACM HTML reader view. For storage, we used Firebase Firestore. For the annotation tooltip, we used the open-source Recogito JS library[1].

4.4 Results

For this exploratory study, our results focus on qualitative findings and observations. Participants experienced both our design probe system Ocean and the standard ACM interface as a grounding point of comparison, and we describe their feedback and performance with both. We observed that most participants found the interface useful overall and utilized different features in Ocean for different tasks. Both low-vision and sighted participants perceived and utilized the system in a similar manner, however, there were more substantial differences between teams.

4.4.1 Creating links for personal benefit. All participants made use of links to digest and navigate concepts in a paper. The positive sentiment toward links was universal among low-vision and sighted readers. For example, a low-vision reader said: “I felt really connected reading. The paper was easy to navigate. It was easy to enter [and come back] to whatever section I’m in.” Users created bookmarks and links as they read to clarify parts that confused them or required clarification. As a sighted reader explained, “Links were helpful to boost what you know, to read the paper and to understand.. [It makes] reading the paper easy.”

Most users also found value in creating links themselves. Out of our 6 users, 5 users (one low-vision, four sighted) said that they took advantage of creating links and that it felt natural to create links during reading. However, one low-vision participant was not used to highlighting or note-taking during reading, so adding links felt to them “like additional tasks to do in the midst of reading”.

---

[1] https://recogito.pelagios.org

Notably, the two groups of participants differed in the kinds of links they created over the week (Table 5). Specifically, Team 1 mostly created links between text referencing core parts of the paper and those parts of the paper, while Team 2 mostly focused on creating links to details about upcoming concepts. While these differences are not statistically significant due to small sample size, participant feedback elaborates on this trend, suggesting that this difference may be due to differences in reading goals.

The majority of links generated by Team 1 were between texts referencing core sections and the corresponding sections. When asked why they focused on adding such links, the team members explained that they assumed the goal of reading and annotating was to prepare for revisiting the paper in the future. For instance, a low-vision team member appreciated the feature as it is challenging for them to skim through a paper when they revisit it: "I like the idea of being able to add anchors on any point of the document for my later reference. That was a nice feature." A sighted team member read the paper as if they were doing "a literature review" and added links for future use. This implies that the participants perceived the link to the structural information in the paper and created their own table of contents, which is known to be useful for low-vision readers [37]. They also expressed a desire to add multiple ends to certain bookmarks. The low-vision user in that team explained: "To be able to [..] chop up the paper and reorganize it in different ways. I just wanted to see everything related to experiment."

On the other hand, links from Team 2 most often connected specific concepts introduced in the paper with text that detailed those concepts. Team members explained their goal in creating these links was to perform better on the quiz tasks. This reported goal aligns with behavior on the quiz and feedback in the exit-survey. Participants tended to do better on the quiz with Ocean, and indicated that navigating information was less mentally demanding because they felt that they had a better understanding of a paper when they used Ocean. They also felt more successful at the task due to their overall improved understanding of the paper, and due to utilization primarily of tabs for finding the required content. Additionally, they had to work less hard to find information that they were looking for when they used Ocean compared to the ACM web interface (Table **??**).

4.4.2 Varied usefulness of navigation features. Participants found different utility in links and tabs for navigating academic papers. As aforementioned, users used links to help unpack information. However, the user-created links were more difficult to use than the pre-populated links when they needed to find specific information, such as during the post-quiz, because of difficulty predicting exactly where the link would take them. A low-vision user said: "When I was reading the paper normally the links were really helpful, but [when] I want to answer a particular question I didn't find [the information that I was looking for]." All the participants initially used links to solve quiz questions, and switched to using search (ctrl+F) when the links did not lead them to the answer. With the limited time frame, either they were not able to tell which link would take them to the desired answer or no relevant links had been created by any readers: "When I had these questions to answer... it was [a] bit finicky to see which link I need to click." Participants (4/6, one low-vision and three sighted) used pre-populated links to solve quiz problems using the Ocean interface. Unlike for user-created links, the destination of pre-populated links to figures, tables, or citations were more clear.

Participants tended to open links in new tabs during the quiz, rather than visiting the end of a link within a current tab. Participants created tabs when they frequently revisited certain parts or compared different parts of a paper. They also used tabs to jump to other parts and to grasp the paper at a high-level. Participants said that the list of tabs they created becomes a "customized table of contents" and it provides them a good summary. A low-vision reader said: "It also gives a kind of a summary of the whole paper, right, like if you can see all the tabs at one point."

4.4.3 Navigation links created by my team vs. crowdsourcing or machine learning. While our participants created links to facilitate their own reading, they also expressed appreciation for links created by their teammates. They used these links to enhance their own paper understanding, and explored some of these links in answering quiz questions. However, relying on a small team to generate all the links for a paper limits scalability, and participants encountered cases where no relevant links had been created by team members to help answer some quiz questions. However, despite the shortcoming of group-generated links, all users said that they would prefer team-sourced links to links created by more scalable solutions like crowdsourcing or machine-generation. Participants said that they valued shared and co-created links with their team members. In particular, they communicated that they would trust links created by their team members more, in terms of usefulness and personal relevance. A low-vision participant said: “I don’t need to know the way other people (crowd workers) interpret the paper.” Using Ocean links, they were able to construct the “shared interpretation through collaborative interaction" which increases the value of the team-sourced links.

The two teams developed distinct collaboration practices using Ocean links. One team said that they wanted to be more hands-on with the links and to engage in close collaboration during link creation. They thought that creating links independently constitutes an incomplete collaboration, and that they would like to deliberate together during the link creation process then tidy links (e.g., make sure that a link from a concept goes to the paragraph that best describes the concept). A low-vision user of the team said: “If there’s no discussion, if there’s no collaboration, for me it’s just random people happening to do something on the same document, but we need to talk in order to... agree upon something before there’s actually collaboration in my mind.” A sighted participant suggested to add an export-links feature in Ocean, so they can put them in a designated shared workspace where team members can go through the links together.

On the other hand, the other team thought that the independently-created links were able to spur more collaboration and enrich their collaboration artifact. Independently creating links allowed them to contribute to the team knowledge with lower efforts, compared to shared comments or annotations with others which required significantly more time and effort. Furthermore, comments and annotations are very personal, whereas links could be more universally helpful. A low-vision reader of the team said, “We know these links connect to here, but in annotations [I think] everybody should have their own way, like they really should read the paper.”

5 DISCUSSION

Our studies highlight and illuminate the major challenges faced by low-vision readers of scientific papers. Discouragingly, half of participants in Study 1 noted that given navigation challenges using current tools, they “gave up” on navigation. Previous researchers have noted that lack of general accessibility has also led to paper abandonment [37]; however, our results show that even for accessible papers, readers may abandon central activities like navigation and instead resort to linear reading. Here, we outline design considerations for academic paper reading interfaces, position of our work relative to hypertext, and discuss limitations and future work.

5.1 Design considerations

From both our studies and previous literature, we consolidated design considerations for academic paper reading interfaces focused on low-vision users, to help address these challenges.

- **Persistent feature access:** Interface components that users can access at any location of the text without scrolling are largely preferable for low-vision users. Examples of such interface components are the fixed header, context menu, and keyboard shortcuts. In both Study

1 and 2, our low-vision participants found interfaces useful that allowed them to activate features without locating buttons or navigating to interface elements, which take a substantial amount of time. There are pros and cons of such perpetual interface components. For a fixed header, the interface is always located at the top, so it is easy for users to predict where to find the interface. However, this is not the case for low-vision users who use system-level magnification. And for browser-magnification users, fixed elements take up valuable screen real estate. Similarly, a context menu allows users access by clicking within the user's current viewport, but overrides the browser's default context menu. Similarly, users can use shortcuts regardless of their magnification setting, but they tax users' cognitive load; users have to remember which key maps to which features. One can imagine a customizable interface that allows users to pick from a suite of interface compositions depending on the settings that a low-vision user is using.

- **Equal access to original text:** During Study 1, low-vision participants wanted to read documents in their original form, even if it enables them to read fewer texts, to understand the same paper, and to make the interface convenient for them. Our results suggest that this may be due to the fear of missing out and misunderstanding the contents. Echoing previous work, separate views for visually impaired and sighted users provides a segregated and typically worse experience for users with disabilities [29]. Low-vision users who read visually can use the same visual interface or presentation as sighted users, and accessibility improvements benefit all users. In line with principles of universal design [29], sighted and low-vision readers should use the same mainstream reading interfaces, which should be fully accessible to all users. Any alterations specifically for low-vision users should be optional and the interface should transparently show any differences from the defaults.
- **Anchors in a sea of text:** Being able to mark or anchor specific texts in a complex document opens up promising opportunities for low-vision readers. In both studies, low-vision participants were able to use such anchors to unpack loaded information. Given that most low-vision readers have to linearly read documents or rely on limited methods of navigation (e.g., table of contents if provided, or ctrl+F), being able to jump through their anchors and quickly go through the text (skim) the document was a huge win. Prior work has shown that it is desirable that anchors could be used in different views and features [12], just as bookmarks and links in Ocean can be used in the main paper view or as tabs. Pre-populated titles for each anchor (e.g., using the parent section title) could be beneficial when anchors are presented in a list to help participants remember link locations.
- **Interfaces that have independence in mind:** Prior work has shown that low-vision technology users often do not want to show their disability or burden others by asking for help [35]. Tools like Ocean that organically crowdsource the required information to build accessible solutions have the potential to help meet these desires. Low-vision participants from Study 2 wanted to forgo automatic approaches (e.g., crowdsourcing and machine learning) for enriching documents, echoing previous work suggesting that introducing automatic technology to visually-impaired users might increase barriers [20]. Furthermore, low-vision users often read papers visually and tend to be more independent than other disability groups [35].
- **Destination clarity of linked content**: Low-vision participants in our studies valued predictability and transparency about where links within the text would take them. In Study 1, participants valued interfaces that showed a preview of linked text without requiring them to make the full jump. And in Study 2, low-vision users made frequent use of pre-populated links in Ocean, since they could accurately predict the linked content. This suggests the need for descriptive cues, content previews, or other clarifications in such interfaces.

- **Provide personal benefits to collaborative accessibility contributors:** A key property of our interface design was that adding bookmarks and links benefits creators themselves (e.g., to understand the paper). In Study 2, our low-vision participants tended to be motivated to create links because it helps them to jump to other parts and clarify information. This suggests that boostrapping data from users [16] may be a promising technique for creating more accessible reading tools. In such a setup, individuals work toward their own reading and navigation goals while also contributing to the overall accessibility of the paper.

5.2 Hypertext for low-vision readers

Our work explores a set of different designs and whether they meet low-vision readers' specific needs. One of the prototypes in Study 1 focuses on navigating information in a paper non-linearly via hyperlinks. Participants found that the interface helped them to clarify information and quickly go through the paper. In Study 2, through the collaborative creation of links, participants indicated that links allowed them to provide or receive guidance from peer readers, and increased the quality of collaboration. These findings resonate with the body of hypertext research; previous research revealed that hypertext allows better navigation and formatting of information sources and is a great means of collaboration [23, 33]. Our work augments the existing body of research by conducting iterative design of reading interfaces focused on low-vision readers and providing empirical results that hypertext interaction could enhance low-vision users' reading experience. In addition, we propose a hypertext-authoring interface that can mitigate the major issues of disorientation and cognitive overhead associated with traditional hyperlinks [6]. In Ocean, users can anchor texts as they read, navigate to other parts in the paper, and only later turn anchors into links when they want to connect, which forgoes premature commitment (allowing for well-thought decisions) [2]. This suggests that the act of adding links is not only embedded in the natural flow of reading, but also may incentivize readers since it allows easy navigation.

In terms of hypertext generation, our work also reveals fruitful insights specific to low-vision users. The general skepticism and technical barriers we identified in Study 1 often cause low-vision users to forgo automated accessibility tools. Team-sourced links are a happy medium compared to crowd- or machine-generated solutions since they help scale the number of links compared to a single author, and generate reliable collaboration artifacts and shared insights between peer researchers. Limiting to only team members might result in a deficiency of links. We encountered pieces of evidence during Study 2 suggesting that creating links requires low effort, and participants were more willing to share links, compared to other forms of traces like annotations, which need more effort and personal information. The mainstream success of hypertext further suggests that the link-based design of Ocean may also benefit sighted users and multi-ability teams.

5.3 Limitations and Future Work

While we provide largely empirical findings and design guidelines based on two user studies, these studies have several limitations. First, both studies had small numbers of participants, largely due to the scarcity of low-vision researchers and resulting recruiting difficulties, which precluded quantitative inferences. In Study 2, we also only deployed our system with two papers picked by the research team. This constraint limits the opportunity for us to see how users might actually use the system in the wild with papers of their own choosing. However, participants engaged in numerous tasks that require thorough reading of the paper, which hopefully resulted in engagement with the papers and our system. Future work should examine how the system can be utilized in the longer term involving more participants. Further studies with larger pools of participants and more design probes are needed to more fully understand user experiences.

Our Ocean prototype also had practical limitations. In particular, our prototype was implemented as a web app, and web readers are less popular than PDF readers for academic papers [37]. However, more and more publishers choose to publish in the HTML format, in large part to expand accessibility, since text in HTML is easily readable and parseable by screen-readers while PDF is often not. For example, SciA11y [37] enables auto-conversion from PDF to HTML. Hence, it is possible that academic readers will increasingly benefit from web readers like Ocean. Furthermore, the design of Ocean can also be applied to PDF and other formats using computer vision and OCR techniques. For example, it may be possible for an interface to support selecting and linking any piece of a PDF document (e.g. a circled area within an image). We built our prototype as a web app based in HTML to get around such unsolved technical problems, and provide our tool with access to pre-parsed content for this initial exploration.

As additional future work, we are interested in further improvements to and applications of reader-generated links. As proposed by one of our participants in Study 2, such links can be useful to other readers outside of a particular research team. For example, deploying a version of Ocean where any paper readers can contribute links would increase scale and access. However, our Study 2 participants warned that links created by such open groups may reduce trust. Figuring out how to boost trust or filter out low-quality link contributions are exciting problems that could be addressed in the future. User feedback from Study 2 might shed light on how to address such challenges. In particular, one user mentioned that they wanted to tidy links based on group discussion. A link-voting interface could help aggregating links when there are many candidate links proposed. Through readers' up-voting and down-voting, teams could decide on which links were optimally pointing to relevant information.

It may also be possible to allow users to customize the look of links and papers via end-user programming interfaces. Previous work suggested the promising implications of end-user programming in document visualization [10]. For example, our Collapse interface in Study 1 was preferred by the participants who wanted to unpack what they read from the paper and go through the same class of information in one place. If the paper contains various Ocean links, the Collapse interface can be implemented as a filter of the links. Especially, since low-vision readers have a broad range of conditions, they could greatly benefit from customized views of academic papers.

6 CONCLUSION

In this work, we explore the challenges that low-vision readers face when navigating academic papers, and explore a design space of academic-paper reading interfaces that organically collect required data from readers. In our first user study involving four new prototype probes, we explored low-vision readers' challenges and front-end interface preferences when reading academic papers; we found that in general, they wanted an interface that would help them orient themselves and digest complex concepts. Next, informed by findings from our first study, we developed a complete system and ran a second study to better understand the end-to-end user experience, including how the required data for such navigation tools might be organically generated by readers. During the study, mixed-ability teams used our prototype reader tool Ocean, which allows users to add and share navigation traces with low effort. Our results suggest that such a tool may have utility to both low-vision and sighted readers navigating academic papers for various goals. Finally, we compiled a list of design considerations for academic paper navigation interfaces for low-vision researchers. We hope that future researchers and developers can leverage these design guidelines and study findings to continue advancing interfaces for low-vision users.

ACKNOWLEDGMENTS

We thank our participants and reviewers for sharing feedback, as well as John Tang and the Microsoft Ability team for insightful guidance.

REFERENCES

[1] Faisal Ahmed, Yevgen Borodin, Andrii Soviak, Muhammad Islam, I.V. Ramakrishnan, and Terri Hedgpeth. 2012. Accessible skimming: faster screen reading of web pages. In *Proceedings of the 25th annual ACM symposium on User interface software and technology (UIST '12)*. Association for Computing Machinery, Cambridge, Massachusetts, USA, 367–378. https://doi.org/10.1145/2380116.2380164

[2] Alan F Blackwell, Carol Britton, A Cox, Thomas RG Green, Corin Gurr, Gada Kadoda, MS Kutar, Martin Loomes, Chrystopher L Nehaniv, Marian Petre, et al. 2001. Cognitive dimensions of notations: Design tools for cognitive technology. In *International Conference on Cognitive Technology*. Springer, 325–341.

[3] Erin Brady, Meredith Ringel Morris, and Jeffrey P. Bigham. 2015. Gauging Receptiveness to Social Microvolunteering. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, Seoul Republic of Korea, 1055–1064. https://doi.org/10.1145/2702123.2702329

[4] Danielle Bragg, Naomi Caselli, John W Gallagher, Miriam Goldberg, Courtney J Oka, and William Thies. 2021. ASL Sea Battle: Gamifying Sign Language Data Collection. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–13.

[5] Stacy M. Branham and Shaun K. Kane. 2015. Collaborative Accessibility: How Blind and Sighted Companions Co-Create Accessible Home Spaces. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 2373–2382. https://doi.org/10.1145/2702123.2702511

[6] Jeff Conklin. 1987. Hypertext: An introduction and survey. *computer* 20, 09 (1987), 17–41.

[7] Maitraye Das, Darren Gergle, and Anne Marie Piper. 2019. "It doesn't win you friends": Understanding Accessibility in Collaborative Writing for People with Vision Impairments. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 1–26. https://doi.org/10.1145/3359293

[8] Nediyana Daskalova, Eindra Kyi, Kevin Ouyang, Arthur Borem, Sally Chen, Sung Hyun Park, Nicole Nugent, and Jeff Huang. 2021. Self-E: Smartphone-Supported Guidance for Customizable Self-Experimentation. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–13.

[9] Paul Dourish, W. Keith Edwards, Anthony LaMarca, and Michael Salisbury. 1999. Presto: an experimental architecture for fluid interactive document spaces. *ACM Transactions on Computer-Human Interaction* 6, 2 (June 1999), 133–161. https://doi.org/10.1145/319091.319099

[10] Philipp Eichmann, Trent Green, Robert Zeleznik, and Andries van Dam. 2017. NuSys: Towards a Document IDE for Knowledge Work. In *Proceedings of the 2017 ACM Symposium on Document Engineering*. 227–230.

[11] Kazuyuki Fujita, Aoi Suzuki, Kazuki Takashima, Kaori Ikematsu, and Yoshifumi Kitamura. 2021. TiltChair: Manipulative Posture Guidance by Actively Inclining the Seat of an Office Chair. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–14.

[12] Gene Golovchinsky and Catherine C Marshall. 2000. Hypertext interaction revisited. In *Proceedings of the eleventh ACM on Hypertext and hypermedia*. 171–179.

[13] Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in psychology*. Vol. 52. Elsevier, 139–183.

[14] Andrew Head, Kyle Lo, Dongyeop Kang, Raymond Fok, Sam Skjonsberg, Daniel S. Weld, and Marti A. Hearst. 2021. Augmenting Scientific Papers with Just-in-Time, Position-Sensitive Definitions of Terms and Symbols. *arXiv:2009.14237 [cs]* (April 2021). http://arxiv.org/abs/2009.14237 arXiv: 2009.14237.

[15] Alan Irwin. 1995. *Citizen science: A study of people, expertise and sustainable development.* Psychology Press.

[16] Ece Kamar and Eric Horvitz. 2012. Incentives for truthful reporting in crowdsourcing.. In *AAMAS*, Vol. 12. 1329–1330.

[17] Steven Komarov and Krzysztof Z Gajos. 2014. Organic peer assessment. In *Proceedings of the CHI 2014 Learning Innovation at Scale workshop*.

[18] Nicholas Kong, Marti A. Hearst, and Maneesh Agrawala. 2014. Extracting references between text and charts via crowdsourcing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Toronto Ontario Canada, 31–40. https://doi.org/10.1145/2556288.2557241

[19] Gordon E. Legge. 2016. Reading Digital with Low Vision. *Visible language* 50, 2 (Aug. 2016), 102–125. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5726769/

[20] A. Lundgard, C. Lee, and A. Satyanarayan. 2019. Sociotechnical Considerations for Accessible Visualization Design. In *2019 IEEE Visualization Conference (VIS)*. 16–20. https://doi.org/10.1109/VISUAL.2019.8933762

[21] Don MacMillan. 2012. Mendeley: teaching scholarly communication and collaboration through social networking. *Library Management* 33, 8/9 (Oct. 2012), 561–569. https://doi.org/10.1108/01435121211279902

[22] Meredith Ringel Morris, Jazette Johnson, Cynthia L. Bennett, and Edward Cutrell. 2018. Rich Representations of Visual Content for Screen Reader Users. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–11. https://doi.org/10.1145/3173574.3173633

[23] Jakob Nielsen. 1990. The Art of Navigating through Hypertext. *Commun. ACM* 33, 3 (mar 1990), 296–310. https://doi.org/10.1145/77481.77483

[24] Allard Oelen, Markus Stocker, and Sören Auer. 2021. Crowdsourcing Scholarly Discourse Annotations. In *26th International Conference on Intelligent User Interfaces*. 464–474.

[25] World Health Organization. 2021. Blindness and vision impairment. https://www.who.int/news-room/fact-sheets/detail/blindness-and-visual-impairment. Accessed 1/10/2022.

[26] Katharina Reinecke and Krzysztof Z Gajos. 2015. LabintheWild: Conducting large-scale online experiments with uncompensated samples. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*. 1364–1378.

[27] Jeffrey M. Rzeszotarski and Meredith Ringel Morris. 2014. Estimating the social costs of friendsourcing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. Association for Computing Machinery, Toronto, Ontario, Canada, 2735–2744. https://doi.org/10.1145/2556288.2557181

[28] Elliot Salisbury, Ece Kamar, and Meredith Morris. 2017. Toward Scalable Social Alt Text: Conversational Crowdsourcing as a Tool for Refining Vision-to-Language Technology for the Blind. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing* 5, 1 (Sept. 2017). https://ojs.aaai.org/index.php/HCOMP/article/view/13301

[29] Anthony Savidis and Constantine Stephanidis. 1995. Developing Dual User Interfaces for Integrating Blind and Sighted Users: The HOMER UIMS. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '95)*. ACM Press/Addison-Wesley Publishing Co., USA, 106–113. https://doi.org/10.1145/223904.223918

[30] David Shotton, Katie Portwin, Graham Klyne, and Alistair Miles. 2009. Adventures in Semantic Publishing: Exemplar Semantic Enhancements of a Research Article. *PLOS Computational Biology* 5, 4 (April 2009), e1000361. https://doi.org/10.1371/journal.pcbi.1000361

[31] Jonathan Silvertown. 2009. A new dawn for citizen science. *Trends in ecology & evolution* 24, 9 (2009), 467–471.

[32] Kevin M. Storer and Stacy M. Branham. 2019. "That's the Way Sighted People Do It": What Blind Parents Can Teach Technology Designers About Co-Reading with Children. In *Proceedings of the 2019 on Designing Interactive Systems Conference (DIS '19)*. Association for Computing Machinery, San Diego, CA, USA, 385–398. https://doi.org/10.1145/3322276.3322374

[33] Norbert Streltz, Frank Halasz, Hiroshi Ishii, Tom Malone, Chris Neuwirth, and Gary Olson. 1991. The role of hypertext for CSCW applications. In *Proceedings of the third annual ACM conference on Hypertext*. 369–377.

[34] Ed Summers, Julianna Langston, Robert Allison, and Jennifer Cowley. 2012. Using SAS/GRAPH® to Create Visualizations That Also Support Tactile and Auditory Interaction. (2012). http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.423.1483&rep=rep1&type=pdf

[35] Sarit Felicia Anais Szpiro, Shafeka Hashash, Yuhang Zhao, and Shiri Azenkot. 2016. How People with Low Vision Access Computing Devices: Understanding Challenges and Opportunities. In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility*. ACM, Reno Nevada USA, 171–180. https://doi.org/10.1145/2982142.2982168

[36] Jayne Wallace, John McCarthy, Peter C. Wright, and Patrick Olivier. 2013. Making Design Probes Work. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) *(CHI '13)*. Association for Computing Machinery, New York, NY, USA, 3441–3450. https://doi.org/10.1145/2470654.2466473

[37] Lucy Lu Wang, Isabel Cachola, Jonathan Bragg, Evie Yu-Yen Cheng, Chelsea Haupt, Matt Latzke, Bailey Kuehl, Madeleine van Zuylen, Linda Wagner, and Daniel S. Weld. 2021. Improving the Accessibility of Scientific Documents: Current State, User Needs, and a System Solution to Enhance Scientific PDF Accessibility for Blind and Low Vision Users. *arXiv:2105.00076 [cs]* (April 2021). http://arxiv.org/abs/2105.00076 arXiv: 2105.00076.

[38] Hui-Yin Wu, Aurélie Calabrèse, and Pierre Kornprobst. 2021. Towards Accessible News Reading Design in Virtual Reality for Low Vision. *Multimedia Tools and Applications* (May 2021). https://hal.inria.fr/hal-03192461

[39] Eran Yogev, Kobi Gal, David Karger, Marc T. Facciotti, and Michele Igo. 2018. Classifying and visualizing students' cognitive engagement in course readings. In *Proceedings of the Fifth Annual ACM Conference on Learning at Scale*. ACM, London United Kingdom, 1–10. https://doi.org/10.1145/3231644.3231648

[40] Dongwook Yoon, Nicholas Chen, Bernie Randles, Amy Cheatle, Corinna E Löckenhoff, Steven J Jackson, Abigail Sellen, and François Guimbretière. 2016. RichReview++ Deployment of a Collaborative Multi-modal Annotation System for Instructor Feedback and Peer Discussion. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. 195–205.

A QUIZ QUESTIONS FOR DASKALOVA ET AL. [8]

What are the things that you can do in the 2nd version of the Self-E interface, but can't do in the initial interface? Select all that apply.

- You can modify the label scale for its "effect".
- You can create your own experiment.
- You can provide two conditions to test out the effect.
- You can customize the daily report and visualization.

How does Self-E help its participants to conduct their self-experiments? Select all that apply.

- By motivating them to change their behavior.
- By helping them realize a new aspect of themselves.
- By helping them readjust the experiment parameters.
- By making them collect more self-awareness data

Select all instances of the previous work "Ayobi et al. [2]" which occur in the paper.

- In the Introduction, presenting the state of previous work on self-experimentation.
- In the Introduction, presenting the lack of evaluation in previous works due to the limited use of domains.
- In the Self-E DESIGN AND IMPLEMENTATION section, motivating their prototype design.
- In the Diary study and findings section, comparing their results with results from previous works.

What is the difference between participants of the first and second study? Select all that apply.

- The second study participants had prior experience of self-experiments.
- The second study participants took a course on experimental design.
- The first study participants were all students.
- The second study participants were a subset of the first study participants who stayed in the study longer than others.

B QUIZ QUESTIONS FOR FUJITA ET AL. [11]

This figure shows the "time limit representation", which is the one of applications of TiltChair. Select all the applications below that are in the same category as the "time limit representation".

- Changing the inclination based on users' non-verbal behavior
- Reminding the user that there is a meeting
- Nudging the user to take a break
- Crunch time

In their evaluation, what did they not use to measure task performance?

- Looking at the number of words participants typed
- Looking at rate between the input and reference texts
- Looking at original NASA-TLX filled in by participants
- Looking at the subjective satisfaction survey filled by participants

The authors conducted the third experiment on a particular application of TiltChair. Select all the applications below that are in the same category as this application.

- Changing the inclination based on users' non-verbal behavior
- Reminding the use that there is a meeting
- Nudging the user to take a break
- Crunch time

Select all instances of the previous work "Bendix et al. [7]" which occur in the paper.

- In the introduction, presenting prior work which passively tilts the seat
- In the related work section, presenting prior work which facilitates trunk movements
- In the TiltChair section, justifying their implementation
- In the User Study section, justifying their study design

C DETAILED QUIZ RESULTS OF STUDY 2

Table 7. Quiz results of each participant: For all participants, their correctness rate either increases or stays same when they use Ocean

Team	**Member vision**	**Completion time (correct rate) using ACM web interface**	**Completion time (correct rate) using Ocean**
1	Low vision	9:22 (0%)	10:00 (25%)
1	Sighted	8:20 (25%)	9:39 (25%)
1	Sighted	7:05 (0%)	5:52 (75%)
2	Low vision	10:00 (0%)	10:00 (0%)
2	Sighted	10:00 (0%)	9:10 (25%)
2	Sighted	6:29 (0%)	3:10 (75%)
Average		8:33 (4.2%)	7:59 (37.5%)
STD		1:30 (10.2%)	2:50 (30.6%)

D COMPLETE LIST OF QUESTIONS OF STUDY 2

Accessibility of the paper

- Did you find the links beneficial for navigating paper? Why or why not?
- Can you compare the interface of two versions – plain HTML and annotated Ocean paper?
- Which links do you think is most helpful for navigation and why?
- What do you think about links authoring interface? Was it natural? (new to the interface? Just complicated)
- Which interface do you prefer for the following tasks: searching for material, skimming, reading as a team, etc.
- What was your motivation of adding links? Do you think the links only help you or your teammates as well? Why and how?

Group dynamics

- What was your motivation of adding links? Do you think the links only help you or your teammates as well? Why and how?
- How did links created by your teammates help your reading?
- How would you feel about links created by crowdworkers or ML model compared to links created by team workers? (i.e. do you think there are links when you needed)
- If the tool were available for future use, would they use it with collaborators? Why or why not?

Received January 2022; revised April 2022; accepted August 2022