

Assessing the Usability of Electronic Theses and Dissertations for Blind Users

Akshay Kolgar Nayak
Old Dominion University
Department of Computer Science
Norfolk, Virginia, USA
anaya001@odu.edu

Hae-Na Lee
Michigan State University
Department of Computer Science and Engineering
East Lansing, Michigan, USA
leehaena@msu.edu

Sampath Jayarathna
Old Dominion University
Department of Computer Science
Norfolk, Virginia, USA
sampath@cs.odu.edu

Vikas Ashok
Old Dominion University
Department of Computer Science
Norfolk, Virginia, USA
vganjigu@odu.edu

Abstract

Electronic Theses and Dissertations (ETDs) provide broad access to student scholarship and support the long-term preservation and dissemination of academic research. Despite their central role in digital libraries, little is known about their usability for people with blindness or visual impairments (BVI) who access scholarly documents using screen readers. We address this gap through a mixed-methods study in which BVI participants completed representative ETD interaction tasks followed by semi-structured interviews. Our findings reveal a substantial gap between accessibility and usability. Although the evaluated ETDs were technically accessible, participants struggled with information-seeking and navigation tasks and experienced considerable cognitive load, even with AI assistance. These challenges reflected limited screen reader AI capabilities, fragmented workflows, and poor integration between screen readers and more capable third-party AI systems. Based on these findings, we derive design considerations for assistive technologies that better support usable interaction with long-form scholarly documents.

CCS Concepts

• **Human-centered computing** → **Accessibility technologies**; **Empirical studies in HCI**; **Usability testing**.

Keywords

Screen reader, Blind, Visually impaired, Digital libraries, User experience

ACM Reference Format:

Akshay Kolgar Nayak, Sampath Jayarathna, Hae-Na Lee, and Vikas Ashok. 2026. Assessing the Usability of Electronic Theses and Dissertations for Blind Users. In *The 2026 ACM/IEEE Joint Conference on Digital Libraries (JCDL '26)*, October 13–16, 2026, Frisco, TX, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3805696.3846043>



This work is licensed under a Creative Commons Attribution 4.0 International License. *JCDL '26, Frisco, TX, USA*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2597-5/2026/10
<https://doi.org/10.1145/3805696.3846043>

1 Introduction

Electronic Theses and Dissertations (ETDs)¹ are a cornerstone of modern digital libraries, providing long-term preservation and broad access to graduate research [11, 54]. ETDs support the discoverability, dissemination, and impact of scholarly work while enabling interdisciplinary knowledge discovery and large-scale analyses such as text mining and scientometrics [6, 14, 21, 42]. Consequently, institutional ETD repositories have become a standard component of digital library infrastructures worldwide [41, 47]. Ensuring that ETDs are not only technically accessible but also practically usable in scholarly tasks is therefore important for equitable participation in research. This is particularly relevant for people with blindness or visual impairments (BVI), who primarily access scholarly documents through screen readers [25, 30, 48].

Although digital accessibility has been extensively studied, existing research has focused primarily on ensuring that digital content meets accessibility requirements such as semantic markup, alternative text, and screen reader compatibility [12, 27, 35]. Less work has examined the *usability* of accessible content, including how efficiently, and with what cognitive effort BVI users complete realistic tasks within ETDs [1, 24, 34, 46]. Existing usability studies have largely examined web browsing, desktop and mobile applications, and general PDF interaction [36, 40, 44]. However, these findings do not directly generalize to ETDs, which are substantially longer than web pages and shorter scholarly documents, exhibit deeper hierarchical structures, and contain chapters, figures, tables, mathematical expressions, citations, cross-references, and indexes that make non-visual navigation and information retrieval more demanding [4, 5, 51]. To the best of our knowledge, no prior work has systematically investigated ETD usability through task-based evaluation coupled with qualitative inquiry.

To address this gap, we conducted a mixed-methods study with 15 BVI participants with prior ETD experience. Participants completed representative tasks involving document navigation and information seeking. The study employed two within-subject conditions: (i) using only a screen reader, and (ii) using a screen reader together with an AI assistant of the participant's choosing. Across

¹In this paper, we use ETD to refer specifically to the thesis or dissertation document deposited in an institutional repository, rather than the repository interface, discovery system, or submission workflow through which it is accessed.

both conditions, we measured task time, task success, screen reader navigation effort, and perceived workload using NASA-TLX. Semi-structured interviews further explored participants' experiences, interaction needs, and design preferences.

Our findings triangulate task performance, interaction logs, workload measures, and participant accounts that highlight how technically accessible ETDs can remain difficult to navigate and interpret non-visually. We highlight how limited page-level orientation, application-specific commands, and cognitively demanding representations of tables and figures contribute to these difficulties, while contemporary AI better supports interpretation than navigation because capable external tools remain disconnected from screen reader and document state. We lastly discuss this gap between accessibility, usability, navigation, and interpretive assistance, and derive design considerations for integrated AI support, document navigation, semantic overviews, and task-oriented interpretation.

2 Related Work

Our work builds on two complementary bodies of literature: (i) accessibility and usability of digital scholarly content for BVI users, and (ii) assistive technologies and AI for non-visual scholarly document interaction.

2.1 Accessibility and Usability of Digital Libraries and Scholarly Content

Digital accessibility research has traditionally focused on making scholarly documents compatible with assistive technologies through semantic markup, alternative text, keyboard operability, and meaningful reading order. The Web Content Accessibility Guidelines (WCAG) provide widely adopted recommendations for making digital content perceivable, understandable, and robust [49]. Related work on PDF accessibility has examined automated tagging, remediation, structural validation, and accessible representations of tables and figures [30, 40]. Universities, publishers, and digital libraries have consequently adopted accessibility requirements and guidance for scholarly PDF submissions, including ETDs [28, 38].

Recent HCI research has increasingly distinguished *technical accessibility* from *practical usability*, arguing that content satisfying accessibility guidelines may still require substantial effort for BVI users to navigate, interpret, and retrieve information [1, 16]. Prior usability studies have shown that screen reader users often perform numerous navigation commands before locating desired information, resulting in increased task completion time, cognitive load, and user frustration [36, 37]. To reduce this burden, researchers have proposed conversational interfaces and LLM-supported information retrieval systems that minimize navigation effort across domains such as web browsing, data visualization, documents, videos, and digital libraries [13, 36, 45, 50].

Accessibility literature in Library and Information Science (LIS) has further shown that BVI users encounter barriers throughout digital-library information seeking, including search formulation, result exploration, resource evaluation, and navigation [39]. Prior work further documents difficulties maintaining awareness of one's location and navigation path, understanding digital-library structures, and determining the format or accessibility of retrieved resources from available metadata [52, 53]. However, this literature

has focused largely on resource discovery and access, leaving less known about usability after a blind reader opens a technically accessible scholarly document.

ETDs provide a particularly useful setting for examining this gap. Their length and hierarchical structure require readers to move across chapters and nested sections and selectively access tables, figures, appendices, and references. These characteristics can amplify screen-reader navigation costs, working-memory demands, and the consequences of losing one's position. Thus, rather than examining barriers at the digital-library search or discovery interface, our study focuses on interaction challenges that persist within a technically accessible scholarly document after retrieval.

2.2 Assistive Technologies and AI for Non-Visual Scholarly Document Interaction

Screen readers remain the primary means by which blind users interact with digital documents and graphical user interfaces [2, 15, 31]. Because they 'read aloud' visually structured content sequentially, users must maintain mental models of document structure and navigate headings, tables, figures, and other elements using application-specific commands [8, 22, 25, 26]. To reduce these interaction demands, recent LLM- and MLLM-based systems have been used to generate image descriptions and alternative text [20, 43], answer questions over documents and charts [36, 48], summarize lengthy content [45], and support natural-language interaction with graphical interfaces [22, 33]. Similar capabilities have also appeared in consumer assistive tools such as Be My AI [7] and Seeing AI [29], enabling users to query visual and document content.

However, prior work largely evaluates isolated AI capabilities or purpose-built interfaces rather than general-purpose AI assistants used alongside screen readers during sustained document interaction. Existing studies also report factual and navigational inaccuracies, reduced performance on long documents, and interaction overhead when AI remains disconnected from screen reader and document state [36, 45, 48]. It therefore remains unclear how such assistants affect navigation and information seeking in long-form scholarly documents such as ETDs. We address this gap in research by examining ETD interaction under screen-reader-only and screen-reader-plus-AI conditions.

3 Participants and Methods

We conducted a mixed-methods, within-subjects study to investigate the usability of technically accessible ETDs for BVI screen reader users and to examine whether contemporary AI assistants improve ETD interaction. The study combined task-based evaluation, interaction logging, workload assessment, and semi-structured interviews. It addressed the following research questions:

- **RQ1:** What usability barriers do BVI screen reader users encounter when performing routine scholarly tasks in ETDs?
- **RQ2:** To what extent do contemporary AI assistants improve the effectiveness, efficiency, and perceived workload of ETD interaction for BVI users?
- **RQ3:** What needs and design preferences do BVI screen reader users have for future assistive technologies supporting interaction with ETDs?

ID	Age/ Gender	Age of Vision Loss	Occupation	Computer Exp. (Years)	Screen Reader(s)	AI Tools
P1	26/F	Since birth	Student	4	NVDA	ChatGPT, Gemini
P2	31/M	Age 9	Accessibility Trainer	8–10	JAWS, NVDA, VoiceOver	ChatGPT, Gemini, Claude
P3	24/M	Don't know	Student	3	JAWS	ChatGPT
P4	21/F	Age 12	Student	5	JAWS, NVDA	ChatGPT
P5	33/M	Since birth	Accessibility Trainer	11	JAWS, NVDA	Claude, ChatGPT
P6	29/F	Age 7	Accessibility Trainer	5	JAWS, NVDA	ChatGPT, Gemini
P7	34/M	Age 20	Self-employed	5–6	JAWS	ChatGPT, Grok
P8	27/F	Age 17	Student	11	JAWS, VoiceOver	Gemini, ChatGPT
P9	30/F	Don't know	Accessibility Trainer	10–11	JAWS, NVDA, VoiceOver	ChatGPT
P10	23/M	Since birth	Student	2–3	NVDA	ChatGPT
P11	32/M	Age 11	Self-employed	7	JAWS, NVDA	ChatGPT, Gemini
P12	22/M	Don't know	Student	5	NVDA	ChatGPT
P13	21/F	Since birth	Student	6	JAWS	ChatGPT
P14	25/F	Age 15	Student	3–4	NVDA	ChatGPT
P15	28/F	Age 22	Self-employed	11	JAWS, NVDA	ChatGPT, Gemini

Table 1: Participant demographics. Participants self-reported all information. All participants were either college students or graduates with prior experience with ETDs.

3.1 Participants

We recruited 15 BVI screen reader users through accessibility-related mailing lists and word-of-mouth referrals. This sample size is comparable to prior in-depth accessibility studies with BVI participants [19, 23] and was intended to support both within-participant comparisons and qualitative investigation of interaction experiences. Participants were eligible if they: (i) were currently enrolled in college or had completed a college degree; (ii) regularly used a desktop screen reader; (iii) had prior experience with scholarly documents including ETDs; (iv) had prior experience using at least one general-purpose AI assistant; and (v) were fluent in English. We excluded individuals whose hearing or motor impairments would prevent them from completing the audio-based computer interaction tasks used in the study.

The sample included 8 women and 7 men between 21 and 34 years of age. Eight participants were university students, four were accessibility trainers, and three were self-employed college graduates. Participants reported between 2 and 11 years of computer experience and used JAWS, NVDA, VoiceOver, or combinations of these screen readers. All participants also reported prior use of AI assistants, most commonly ChatGPT and Gemini. Table 1 summarizes their characteristics.

3.2 ETD Materials and Study Tasks

We selected two ETDs from the same institutional repository.² The documents represented different disciplinary contexts but were selected to be broadly comparable in length and structural complexity and contained common ETD elements, including a table of contents, chapters and subsections, figures, tables, references, and

other navigational structures. Our selection therefore prioritized comparable document structure rather than disciplinary similarity. To prevent technical accessibility violations from dominating the observed interaction difficulties, we remediated both PDFs using Adobe Acrobat Pro. This allowed us to examine whether structurally accessible ETDs remained practically usable during realistic interaction tasks.

Participants completed four representative ETD tasks:

- **T1—Document overview:** Locate the table of contents to understand how the thesis is organized.
- **T2—Hierarchical navigation:** Reach a named subsection within a specified chapter to access a known part of the thesis.
- **T3—Table-based information seeking:** Starting from the page containing the target table, answer a question requiring information to be identified and compared across table entries.
- **T4—Figure-based information seeking:** Starting from the page containing the target figure, answer a question requiring interpretation of information presented in the figure.

We selected four representative tasks spanning two interaction goals relevant to ETD use: document navigation (T1–T2) and local information seeking and interpretation (T3–T4). These tasks were intended to capture specific usability challenges rather than exhaustively represent all ETD activities. Parallel task variants were constructed for the two ETDs so that participants performed the same types of activities without repeating an identical task on the same document.

3.3 Study Conditions

Each participant completed the tasks under two conditions:

²<https://deepblue.lib.umich.edu/items/1d98b5c2-a3c8-4921-8d76-a76e7b4da671> and <https://deepblue.lib.umich.edu/items/fb228b59-698a-4056-a922-ecfcc537364e>.

- **Screen Reader Only (SRO):** Participants completed the tasks using their preferred screen reader without assistance from external AI tools.
- **Screen Reader + AI (SAI):** Participants could use their preferred general-purpose AI assistant in conjunction with their screen reader. Participants determined when and how to consult the AI assistant, thereby reflecting their established interaction practices.

We used a within-subjects design because screen reader proficiency and navigation strategies vary substantially across individuals. Having each participant complete both conditions, therefore, enabled direct comparison while controlling for individual differences. To mitigate learning, document, and order effects, we counterbalanced both condition order and ETD assignment using a Latin-square design [9], such that both ETDs appeared across the SRO and SAI conditions and neither condition systematically occurred first. Participants were distributed as evenly as possible across the resulting sequences. At the beginning of every condition, the screen reader focus was placed at the top of the ETD's first page for navigational tasks (T1 and T2) and the top of the page containing the table or figure for the information-seeking tasks (T3 and T4). This starting-point design ensured separation of document-level navigation demands in T1 and T2 from local content interpretation demands in T3 and T4. Participants used the same computer configuration and screen reader settings across their two conditions.

3.4 Measures

We collected complementary measures of effectiveness, efficiency, interaction effort, and perceived workload. For every task in both conditions, we recorded:

- **Task success**, coded according to whether the participant reached the requested location or provided the correct requested information;
- **Task time**, measured from presentation of the task prompt until completion or expiration of the time limit; and
- **Screen reader navigation effort**, operationalized as the number of screen reader shortcut presses used during the assigned task.

We also assessed perceived workload under each condition, using the NASA-TLX questionnaire [17]. Each task had a maximum duration of 10 minutes. Tasks not completed correctly within this period were coded as unsuccessful. The experimenter also documented notable navigation strategies, breakdowns, errors, and recovery behaviors to contextualize the quantitative measures.

3.5 Interview Design

Following the task-based portion of the study, participants completed a semi-structured interview examining their experiences with ETDs and AI-assisted scholarly document interaction. We developed the protocol using established interviewing guidelines [32] and prior studies involving BVI participants [3, 24]. The interview addressed participants' prior ETD practices, challenges experienced during the study, perceptions of the two conditions, unmet interaction needs, and preferences for future assistive technologies. Example seed questions included:

- *"What challenges do you typically encounter when interacting with ETDs, and which of these challenges did you experience during today's tasks?"*
- *"How did the availability of an AI assistant affect the way you approached the tasks?"*
- *"In which situations did the AI assistant help, and in which situations was it inaccurate, inefficient, or difficult to use?"*
- *"What forms of coordination between a screen reader and an AI assistant would make ETD interaction more usable?"*
- *"What concerns would you have about relying on AI for navigating or interpreting scholarly documents?"*

The protocol was semi-structured rather than fixed. The interviewer used follow-up questions to clarify participants' statements, elicit concrete examples, probe unexpected experiences, and explore apparent inconsistencies [18].

3.6 Study Procedure

The study protocol was reviewed and approved by our institution's Institutional Review Board (IRB). Before beginning the study, each participant reviewed an accessible consent form and provided informed consent. Participants received a \$100 gift card upon completing the session. Each session began with a demographic questionnaire and a brief discussion of the participant's prior experience with ETDs, screen readers, and AI assistants. The experimenter then introduced the task procedure and confirmed that the participant could operate the study computer using their preferred screen reader. Participants completed the two conditions in their assigned counterbalanced order, with the corresponding ETD loaded at the beginning of each condition. The experimenter read each task prompt aloud and provided it in an accessible digital format but did not offer assistance with completing the task. After each condition, participants completed the NASA-TLX assessment for the tasks performed in that condition. The semi-structured interview followed completion of both conditions and lasted approximately one hour. With participants' permission, the interviews were audio-recorded and subsequently transcribed for analysis.

3.7 Quantitative Data Analysis

We analyzed the quantitative data at both the individual-task and participant levels. We first computed descriptive statistics for task success, completion time, screen reader shortcut presses, and NASA-TLX ratings. Because the study used a within-subjects design and included a relatively small sample, we used paired non-parametric tests rather than assuming normally distributed measurements. For continuous and ordinal outcomes, including completion time, shortcut presses, and NASA-TLX ratings, we used Wilcoxon signed-rank tests for paired comparisons. Task-success rates were compared using paired binary analyses, e.g., McNemar tests. We compared the SRO and SAI conditions for each of the tasks, reflecting their distinct interaction demands. Because the study involved multiple task-level comparisons, we applied Holm correction separately within each family of related statistical tests (i.e., task time, screen reader command count, and task success) to control the family-wise error rate. We report Holm-adjusted p -values together with the corresponding effect sizes. Unless otherwise specified, statistical tests used a two-sided significance threshold of $\alpha = .05$.

3.8 Qualitative Data Analysis

We analyzed the interview transcripts and observational notes using reflexive thematic analysis following Braun and Clarke's six-phase process [10]. This approach was appropriate because our goal was to develop an interpretive understanding of participants' ETD interaction experiences, usability barriers, AI-related challenges, and design preferences rather than to quantify predefined categories. The analysis combined inductive coding grounded in participants' accounts with deductive interpretation informed by the study's research questions.

The first author served as the primary analyst. The first author and a second researcher independently conducted iterative, line-by-line coding of assigned transcripts after familiarizing themselves with the data, using in-vivo codes where participants' own wording captured important ideas. The first author subsequently reviewed and synthesized the complete set of codes across all transcripts, developing candidate themes that were further refined through collaborative discussions with the research team. Throughout this process, observational notes were used to contextualize participants' interview accounts and examine whether reported experiences aligned with behaviors observed during the study tasks. Team discussions considered alternative interpretations and divergent cases before refining and naming the final themes. Consistent with reflexive thematic analysis, these discussions were intended to enrich interpretation rather than establish coding reliability or inter-coder agreement [10].

4 Results

We report medians and interquartile ranges (IQRs) because the study involved paired observations from a relatively small sample and several measures exhibited ceiling effects or skewed distributions. We compared SRO and SAI using two-sided Wilcoxon signed-rank tests for task time, screen reader command count, and NASA-TLX workload. Rank-biserial correlation (r_{rb}) is reported as the effect size, with larger absolute values indicating stronger condition differences. Task success was compared using exact McNemar tests. Holm correction was applied separately across the four task-level comparisons for each outcome family. Unless otherwise stated, reported p -values are Holm-adjusted.

4.1 T1: Navigating to the Table of Contents

AI assistance did not significantly improve performance on T1. Median task time decreased from 459 seconds in SRO (IQR: 352.5–561) to 402 seconds in SAI (IQR: 307–551.5), but the difference was not statistically significant ($W = 31$, $p = .530$, $|r_{rb}| = .21$) (Figure 1a). Screen reader command counts were also comparable between SRO ($Mdn = 239$, IQR: 134–291) and SAI ($Mdn = 154$, IQR: 118–356; $W = 50$, $p = 1.000$, $|r_{rb}| = .17$) (Figure 1b). Eleven participants completed the task successfully in SRO, compared with 12 in SAI; this difference was not significant (exact McNemar $p = 1.000$) (see Figure 1c).

Behavioral observations help explain the limited effect of AI. In SRO, participants primarily relied on basic arrow- and tab-based

navigation rather than more efficient PDF-specific commands. Participants sometimes confused web-navigation commands with PDF-reader commands, causing unexpected focus changes, disorientation, and, in some cases, a restart from the beginning of the document. Several participants also attempted to use the PDF reader's search function, but broad queries such as "table of contents" returned multiple matches and provided limited support for reaching the desired heading.

AI did not remove these navigation barriers. Some participants did not consult AI because they did not expect it to control or assist navigation within the PDF reader. Others uploaded the ETD and asked for the location of the table of contents. Although AI could often provide a page number or reproduce the table of contents in its own interface, participants still had to return to the PDF reader and navigate manually. A few participants requested sequences of screen reader commands from AI and alternated between the AI and PDF interfaces while executing the instructions. These instructions were sometimes inconsistent with the active PDF interface or screen reader state, introducing additional errors and recovery effort. Thus, AI provided information *about* the document but could not directly execute or reliably guide navigation within it.

4.2 T2: Navigating to a Specified Subsection

T2 was the most demanding navigation task under both conditions. Median task time was 522 seconds in SRO (IQR: 472.5–600) and 445 seconds in SAI (IQR: 394–600). Although the reduction favored SAI and had a moderate-to-large effect size, it was not statistically significant after correction ($W = 12$, $p = .228$, $|r_{rb}| = .56$) (Figure 2a). Similarly, the median screen reader command count decreased from 296 in SRO (IQR: 253.5–350) to 287 in SAI (IQR: 164.5–319), but the difference was not significant ($W = 43$, $p = 1.000$, $|r_{rb}| = .28$) (Figure 2b). Eight participants completed T2 successfully in SRO, compared with 10 in SAI; this difference was also not significant (exact McNemar $p = 1.000$) (Figure 2c).

Participants exhibited many of the same navigation strategies and breakdowns observed in T1, but the deeper document hierarchy increased their impact. Locating a specified subsection required participants to maintain awareness of both chapter-level and subsection-level structure while traversing a long sequential representation. Participants frequently reverted to basic navigation commands when they could not recall the relevant PDF-specific shortcuts. Experimenter observations indicated more frequent pauses, requests for breaks, and verbal expressions of fatigue than in T1.

As in T1, AI could sometimes identify the likely chapter, subsection, or page number, but it could not transfer screen reader focus to that location. Participants, therefore, remained responsible for translating the AI-provided information into application-specific navigation actions. Switching between the AI and PDF interfaces also disrupted their awareness of the current focus position. Overall, SAI showed a numerical advantage, but AI did not reliably overcome the interaction costs associated with hierarchical navigation.

4.3 T3: Interpreting Information in a Table

AI assistance produced substantial improvements on T3. Median task time decreased from 258 seconds in SRO (IQR: 196–311.5) to 154 seconds in SAI (IQR: 96–192) (Figure 3a). Every participant




Figure 1: Task T1 results.




Figure 2: Task T2 results.

completed the task more quickly in SAI than in SRO, yielding a significant difference with the maximum observed effect size ($W = 0$, $p = .002$, $|r_{tb}| = 1.00$). The median screen reader command count similarly decreased from 125 in SRO (IQR: 101–145.5) to 74 in SAI (IQR: 58.5–86; $W = 0$, $p = .003$, $|r_{tb}| = 1.00$) (Figure 3b).

Task success increased from 9 of 15 participants (60%) in SRO to 15 of 15 participants (100%) in SAI. The unadjusted exact McNemar test was significant ($p = .031$), but the comparison did not remain significant after Holm correction across the four task-success tests ($p = .125$) (Figure 3c). The improvement nevertheless represents six participants who answered incorrectly in SRO but successfully completed the task in SAI, with no changes in the opposite direction.

In SRO, participants navigated repeatedly among table cells while attempting to retain and compare multiple numerical values. Several participants lost track of the relationship between cell values and their row or column headers, producing incorrect comparisons despite the table being technically accessible. Participants

also repeatedly revisited cells because they had forgotten previously encountered values, increasing both navigation effort and cognitive demand.

In SAI, all participants sought AI assistance at the beginning of the task. Participants generally captured an image of the page containing the table and provided it to an AI assistant. This strategy reduced the need for sequential cell-by-cell navigation and enabled AI to perform the requested comparison directly. However, the process was not frictionless. Some screenshots omitted part of the table, requiring participants to repeat the capture-and-upload process. Several participants also returned to the PDF to manually verify the AI-generated answer before reporting it. Despite these additional steps, SAI substantially reduced task time and screen reader command use and enabled all participants to provide the correct answer.




Figure 3: Task T3 results.




Figure 4: Task T4 results.

4.4 T4: Interpreting Information in a Figure

AI assistance also significantly improved figure interpretation. Median task time decreased from 292 seconds in SRO (IQR: 270–321.5) to 225 seconds in SAI (IQR: 208.5–240; $W = 4$, $p = .007$, $|r_{rb}| = .92$) (Figure 4a). Unlike T3, however, the number of screen reader commands did not differ significantly. Participants used a median of 33 commands in SRO (IQR: 25–41) and 36 commands in SAI (IQR: 25.5–41; $W = 49.5$, $p = 1.000$, $|r_{rb}| = .18$) (Figure 4b).

As in T3, task success increased from 9 of 15 participants (60%) in SRO to 15 of 15 participants (100%) in SAI (Figure 4c). The unadjusted exact McNemar test was significant ($p = .031$), but the result did not remain significant after Holm correction ($p = .125$). All six discordant cases favored SAI.

In SRO, participants typically listened to the figure's alt-text multiple times before answering. Although the alt-text contained the information needed to complete the task, participants had to retain multiple labels, values, and trends in working memory while

determining which label corresponded to the highest value. Participants also listened to nearby text to establish context and frequently expressed listening and mental fatigue during repeated playback.

In SAI, all participants used AI assistance from the beginning. Many initially used AI functionality embedded in their screen reader before consulting a third-party assistant. Participants followed two primary strategies: providing only the figure to the AI or capturing both the figure and its surrounding text. The latter strategy supplied additional scholarly context that was not available when the figure was processed in isolation. Participants often asked the AI for confirmation or reformulated their questions before finalizing an answer, but they rarely returned to the PDF after obtaining an adequate response. These behaviors suggest that AI accelerated the extraction and interpretation of figure content, even though interacting with the AI interface required a comparable number of screen reader commands.

4.5 Perceived Workload

Participants reported lower workload in the AI-assisted condition. The median NASA-TLX score decreased from 66.67 in SRO (IQR: 62.33–67.67) to 51.67 in SAI (IQR: 48.83–59.33). All 15 participants reported lower workload in SAI than in SRO, producing a statistically significant difference with the maximum observed effect size ($W = 0, p < .001, |r_{tb}| = 1.00$).

This reduction should be interpreted alongside the task-specific findings. AI did not significantly improve the two document-navigation tasks, but it substantially reduced the effort required to interpret tables and figures. Participants' overall workload ratings therefore appear to reflect the value of replacing repeated sequential inspection and working-memory-intensive comparison with direct natural-language assistance, even though interface switching and manual navigation remained burdensome.

4.6 Qualitative Findings

Reflexive thematic analysis of the interviews identified four themes that contextualize and extend the task-level findings.

4.6.1 Alternative Text Provides Access but Imposes Substantial Interpretive Effort. Participants consistently distinguished the presence of alternative text from the usability of that text. Although detailed descriptions made charts and figures technically accessible, participants found it difficult to retain multiple values, compare trends, and construct an overall understanding through sequential audio alone. Repeated listening increased both cognitive and auditory fatigue. P6 explained:

“Descriptive text for graphs is tiring. Too much to keep in mind as you listen. Then mentally comparing values and figuring out what the graph is conveying is overwhelming to me most of the time.”

Participants therefore wanted representations that communicated relationships and salient trends rather than merely verbalizing visual elements. Some suggested dialog-based presentations in which multiple speakers progressively discussed the important findings in a figure. P6 continued:

“I would rather the screen reader provide information as a dialog between two parties. As the dialog flows, the parties can proactively discuss the salient trends in data. This will also be very engaging and less stressful.”

These accounts reinforce the T4 results: even sufficiently detailed alternative text required substantial repeated listening and mental comparison, whereas AI could more directly answer a focused question about the figure.

4.6.2 Efficient Navigation Requires Better Page-Level Orientation. Participants described page-level navigation as a persistent weakness of current PDF interaction. Many knew that PDF readers offered thumbnail-based navigation but could not recall the application-specific commands required to access it. This difficulty often caused them to fall back on slow, linear navigation even when a more efficient feature existed.

Participants also noted that conventional thumbnail interfaces provide little useful nonvisual information about page content. They wanted dynamically generated summaries that would communicate

the main contents of each page while moving through thumbnails. P9 stated:

“When I go from page to page searching for something, I would like to hear a summary of what is there in that page instead of just the first line in that page. This will help me decide quickly if I have to jump to the next page or search closer in the current page.”

Such summaries could support rapid orientation without requiring users to enter each page and inspect its contents sequentially. This need reflects the difficulties observed in T1 and T2, where technical document structure did not provide an efficient overview of the ETD.

4.6.3 Application-Specific Commands Fragment Screen Reader Interaction. Participants found it difficult to remember and distinguish the different screen reader commands required across web browsers, PDF readers, and other applications. Even experienced users reported confusing command sets or forgetting efficient application-specific combinations. When this occurred, they reverted to generic arrow- and tab-based navigation, substantially increasing interaction time and effort in long ETDs.

Participants wanted customizable or reusable commands that worked consistently across applications, potentially enabled through AI-mediated translation of familiar commands into application-specific actions. P2 explained:

“It will be great if I can just use the web shortcuts I am very familiar with in theses too. Can AI help me with that? I think such a feature will dramatically increase my speed and productivity in theses documents.”

This theme provides qualitative support for the navigation-task results. AI assistants could describe where content was located or suggest commands, but they could not normalize the interaction vocabulary across interfaces or execute the corresponding screen reader-navigation actions.

4.6.4 AI Support Is Divided Between Limited Embedded Features and Powerful but Disconnected External Tools. Participants valued the increasing availability of AI within screen readers, particularly for obtaining descriptions of figures and charts. However, they considered current embedded functionality too narrow for scholarly reading. Participants reported that image descriptions often lacked awareness of surrounding text, introduced unexplained abbreviations, or failed to connect a figure with relevant discussion elsewhere in the ETD. P14 stated:

“I don't have to take a snapshot every time to get both text and image. If I ask for an image description, the AI should be smart enough to look at related text before generating the description.”

Participants also wanted embedded AI to assist with textual comprehension and cross-document navigation rather than only visual description. P11 explained:

“I just don't need help with figures. I need help with the text too. If I don't understand something, I would like AI to help me understand that part or redirect me to another chapter in the thesis where I can get more contextual details.”

Third-party AI assistants offered broader question-answering and interpretation capabilities but remained disconnected from the screen reader and PDF application. Participants described taking screenshots, locating and uploading files, and repeatedly switching interfaces as time-consuming and frustrating. P12 summarized the desired interaction:

“I wish that it took just three or four shortcuts for me to select what I want in a thesis and ask questions about it using my favorite ChatGPT. I don’t want to go through hoops to do this.”

Together, these findings reveal a division between embedded AI that is integrated but limited and third-party AI that is more capable but operationally disconnected. Participants wanted an integrated assistant that could access the current document context, answer questions about both text and visual content, and directly coordinate with screen reader navigation rather than requiring manual transfer of information across applications.

5 Discussion

Our findings reinforce the distinction between technical accessibility and practical usability established in prior accessibility research [1, 16]. Prior work on accessible PDFs has largely focused on ensuring structural access through tagging, reading-order correction, alternative text, and remediation [30, 40]. We extend this work by showing that these provisions do not necessarily make long-form scholarly documents efficient to navigate or interpret. More importantly, our findings reveal why these difficulties persist and where contemporary AI does and does not reduce them. In the remainder of this section, we revisit these findings in light of our research questions and then derive design guidelines for AI-Powered Assistive Technologies that enable usable access to ETDs for BVI users.

5.1 Answering the Research Questions

RQ1: What usability barriers do BVI screen reader users encounter while interacting with ETDs? The primary barriers arose from the interaction between long document structures and sequential screen reader access. Prior work has shown that screen reader users construct mental models by sequentially traversing interface elements and rely heavily on structural cues and learned navigation strategies [8, 26]. ETDs amplify these demands: headings, figures, tables, and cross-references may be technically exposed, but locating and relating them across a long hierarchy requires users to repeatedly establish where they are and what surrounds the current element. This helps explain why remediation alone did not eliminate navigation effort in our study. Similarly, detailed alternative text provided access to figures but required participants to retain and compare multiple values through sequential audio. Thus, the remaining usability cost was not primarily missing information, but the effort required to orient, integrate, and act on information that was technically available.

RQ2: To what extent do contemporary AI assistants improve ETD interaction? AI produced markedly different effects for interpretation and navigation. For tables and figures, AI could compress information that otherwise required repeated sequential inspection into a direct, task-relevant response. This reduced

the working-memory and comparison demands associated with listening to multiple values, consistent with prior work showing the value of conversational access to complex or distributed information [13, 36, 48]. Navigation required a different capability. Reaching a location depended on the user’s current focus, document hierarchy, PDF-reader state, and application-specific screen reader commands. External AI assistants had access to document content but not this interaction state, while embedded screen reader AI had access to the interface but offered more limited reasoning capabilities. This separation explains why AI improved content interpretation but produced little benefit for navigation. The limitation therefore appears to lie not only in model capability, but in the lack of coordination between document understanding and assistive interaction.

RQ3: What needs and design preferences do BVI users have for future assistive technologies? Participants’ preferences further suggest that improving AI support requires tighter coupling with existing screen reader workflows rather than simply adding more capable question answering. Prior work has explored LLMs as a means of reducing interaction complexity and translating user intent into interface actions [22]. Our findings extend this direction to scholarly documents: participants wanted assistance that knew their current position, could interpret surrounding document context, and could act through familiar screen reader interactions. Such integration could support both orientation and interpretation without requiring users to repeatedly move between applications or reconstruct context. This motivates assistive systems in which AI operates as part of the navigation loop rather than as a separate destination for document questions.

5.2 Design Considerations for Future AI-Powered Assistive Technologies

Our findings suggest several opportunities for improving non-visual interaction with scholarly documents by reducing interface fragmentation, supporting efficient navigation, and providing context-sensitive assistance.

DC1: Integrate AI directly into screen reader workflows. Rather than requiring users to switch between PDF readers and external AI assistants, future systems should provide seamless AI assistance within the screen reader itself. Such integration would reduce interaction overhead while allowing AI to access the current document, page, selected element, and navigation state. This would enable users to ask context-dependent questions without repeatedly uploading files, capturing screenshots, or re-establishing their location in the document.

DC2: Support AI-assisted document navigation. Current AI assistants effectively answer questions about document content but provide limited support for reaching the relevant content. Future systems should enable AI to interpret document structure, identify target chapters or sections, and coordinate directly with screen-reader and PDF-reader commands to move users to relevant locations. The system should also confirm the destination and preserve the user’s prior position so that navigation actions remain understandable and reversible.

DC3: Provide semantic document overviews. Long ETDs require users to repeatedly inspect pages to determine whether

they contain relevant information. AI-generated page summaries, chapter overviews, and semantic thumbnail descriptions could substantially improve orientation and reduce navigation effort during exploratory reading. These summaries should be available at multiple levels of detail and linked to the corresponding source locations so that users can move directly from an overview to the relevant content in the ETD.

DC4: Move beyond descriptive accessibility toward interpretive assistance. Participants valued AI not simply because it verbalized inaccessible content, but because it could transform dense, sequential information into task-relevant explanations. Future assistive systems should therefore support interpretive assistance, helping users identify relationships, compare values, summarize trends, and answer questions while remaining grounded in the source. For example, rather than reading every bar in a chart, a system could explain which category changed most and why that pattern is notable; in a large table, it could identify outliers or compare selected rows; and in a long policy or technical document, it could synthesize related requirements scattered across sections. Such assistance should remain inspectable by identifying the source elements underlying each interpretation and allowing users to return directly to them for verification.

5.3 Limitations

This study has a few limitations that should be considered when interpreting the findings. First, the study included 15 BVI participants, which is comparable to prior accessibility research involving blind users but nevertheless limits statistical power and the generalizability of the findings. Larger-scale studies involving more diverse participant populations would strengthen external validity. Second, although the study used two ETDs from different disciplinary contexts, they cannot represent the diversity of ETDs available across disciplines. We also did not control for alignment between participants' academic backgrounds and the ETD subject matter, which may have influenced content-intensive tasks. Future work should investigate documents from a broader range of subject areas and examine whether document domain influences navigation and information-seeking performance.

Third, both ETDs were remediated to satisfy accessibility requirements, representing a comparatively favorable case rather than the range of unremediated or structurally complex ETDs encountered in practice. Future work should therefore examine ETDs with varying levels of accessibility and structural complexity. Fourth, the study was conducted in a controlled laboratory setting using four predefined tasks that captured selected forms of navigation and information seeking rather than the full range of ETD activities, such as exploratory reading, cross-document comparison, citation tracing, or sustained research workflows. Future longitudinal studies should examine AI-supported ETD interaction in participants' everyday academic workflows. Finally, while all participants had prior experience reading ETDs, the sample consisted primarily of university students and college-educated adults. Future work should examine broader populations, including lifelong learners and members of the general public who access ETDs through institutional and public repositories.

Beyond these limitations, our study intentionally evaluated participants' existing AI-assisted workflows rather than a standardized AI system. Participants used different AI assistants, including ChatGPT, Gemini, Claude, and Grok, which vary in long-document understanding and multimodal processing capabilities. Participants also differed in the number of AI tools they used, introducing additional variability in familiarity and tool-selection strategies. Although this design reflects current real-world practice, these differences may have influenced task performance. Future work should compare specific AI models under controlled conditions and standardize tool exposure to better isolate the contribution of AI capabilities to scholarly document accessibility.

6 Conclusion

Electronic Theses and Dissertations (ETDs) are among the valuable resources in digital libraries, making their effective use essential for equitable participation in scholarly research. Through a mixed-methods study with 15 BVI users, we showed that technical accessibility alone does not guarantee practical usability. Although the evaluated ETDs satisfied accessibility requirements, participants experienced substantial challenges navigating long document structures and interpreting complex visual content. Contemporary AI assistants significantly improved table and figure interpretation, reducing task time and perceived workload, but offered limited support for document navigation because they remained disconnected from screen reader workflows. Building on these findings, we identified key usability barriers, uncovered users' needs and preferences, and proposed design considerations for future AI-powered assistive technologies that tightly integrate document understanding, navigation, and screen reader interaction.

Acknowledgments

This work was supported by the Institute of Museum and Library Services (IMLS) through the National Leadership Grants for Libraries program (Award No. LG-256693-OLS-24). The views, findings, conclusions or recommendations expressed in this paper do not necessarily represent those of the Institute of Museum and Library Services.

References

- [1] Amaia Aizpurua, Simon Harper, and Markel Vigo. 2016. Exploring the relationship between web accessibility and user experience. *International Journal of Human-Computer Studies* 91 (2016), 13–23.
- [2] Apple Inc. 2022. Accessibility - Vision - Apple. <https://www.apple.com/accessibility/vision/>.
- [3] Catherine M. Baker, Lauren R. Milne, and Richard E. Ladner. 2019. Understanding the Impact of TVIs on Technology Use and Selection by Children with Visual Impairments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, Article 424, 13 pages. doi:10.1145/3290605.3300654
- [4] Bipasha Banerjee, William A. Ingram, and Edward A. Fox. 2024. Automating chapter-level classification for electronic theses and dissertations. In *2024 IEEE International Conference on Big Data (BigData)*. IEEE, 2400–2409.
- [5] Bipasha Banerjee, William A. Ingram, Jian Wu, and Edward A. Fox. 2022. Applications of data analysis on scholarly long documents. In *2022 IEEE International Conference on Big Data (Big Data)*. IEEE, 2473–2481.
- [6] Siviwe Bangani. 2018. The impact of electronic theses and dissertations: a study of the institutional repository of a university in South Africa. *Scientometrics* 115, 1 (2018), 131–151.
- [7] Be My Eyes. 2026. Be My AI. <https://www.bemyeyes.com/bme-ai/>.

- [8] Yevgen Borodin, Jeffrey P Bigham, Glenn Dausch, and IV Ramakrishnan. 2010. More than meets the eye: a survey of screen-reader browsing strategies. In *Proceedings of the 2010 international cross disciplinary conference on web accessibility (W4A)*. 1–10.
- [9] James V Bradley. 1958. Complete counterbalancing of immediate sequential effects in a Latin square design. *J. Amer. Statist. Assoc.* 53, 282 (1958), 525–528.
- [10] Virginia Braun and Victoria Clarke. 2021. *Thematic Analysis: A Practical Guide*. SAGE Publications Ltd.
- [11] Michael Connolly, Christof Lutteroth, and Beryl Plimmer. 2010. Document resizing for visually impaired students. In *Proceedings of the 22nd Conference of the Computer-Human Interaction Special Interest Group of Australia on Computer-Human Interaction*. 128–135.
- [12] World Wide Web Consortium et al. 2008. Web content accessibility guidelines (WCAG) 2.0. (2008).
- [13] Andrea Esposito, Rosa Lanzilotti, Maristella Matera, Antonio Piccinno, Ludovica Piro, Emanuele Pucci, and Grazia Ragone. 2026. Conversational AI for Digital Accessibility: An Experimental Study Involving Blind and Low Vision Users. *International Journal of Human-Computer Interaction* (2026), 1–24.
- [14] Tránsito Ferreras-Fernández, Francisco García-Peñalvo, José A Merlo-Vega, and Helena Martín-Rodero. 2016. Providing open access to PhD theses: visibility and citation benefits. *Program* 50, 4 (2016), 399–416.
- [15] Freedom Scientific. 2022. JAWS® – Freedom Scientific. <https://www.freedomscientific.com/products/software/jaws/>.
- [16] Stéphanie Giraud, Pierre Thérouanne, and Dirk D Steiner. 2018. Web accessibility: Filtering redundant and irrelevant information improves website usability for blind users. *International Journal of Human-Computer Studies* 111 (2018), 23–35.
- [17] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Human Mental Workload*, Peter A. Hancock and Najmedin Meshkati (Eds.). Advances in Psychology, Vol. 52. North-Holland, 139–183. doi:10.1016/S0166-4115(08)62386-9
- [18] Siw Elisabeth Hove and Bente Anda. 2005. Experiences from Conducting Semi-Structured Interviews in Empirical Software Engineering Research. In *Proceedings of the 11th IEEE International Software Metrics Symposium (METRICS 2005)*. IEEE. doi:10.1109/METRICS.2005.24
- [19] Earl W. Huff Jr, Kwajo Boateng, Makayla Moster, Paige Rodeghero, and Julian Brinkley. 2021. Exploring the Perspectives of Teachers of the Visually Impaired Regarding Accessible K–12 Computing Education. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education (SIGCSE '21)*. Association for Computing Machinery, 156–162. doi:10.1145/3408877.3432418
- [20] Gia Ky Huynh and Wenjun Lin. 2024. SmartCaption AI-Enhancing Web Accessibility with Context-Aware Image Descriptions Using Large Language Models. In *2024 International Conference on Computer and Applications (ICCA)*. IEEE, 1–7.
- [21] Masana Hlengiwe Michelle Khanyi. 2026. *Insights into the South African research landscape through mining theses and dissertations using transformer-based language models*. Ph. D. Dissertation. Stellenbosch: Stellenbosch University.
- [22] SR Kodandaram, U Uckun, X Bi, I Ramakrishnan, and V Ashok. 2024. Enabling uniform computer interaction experience for blind users through large language models [Paper presentation]. In *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility*, St. John's, NL, Canada. <https://doi.org/10.1145/3663548.3675605>.
- [23] Akshay Kolgar Nayak, Yash Prakash, Md Javedul Ferdous, Sampath Jayarathna, Hae-Na Lee, and Vikas Ashok. 2026. Examining Inclusive Computing Education for Blind Students in India. In *Proceedings of the 57th ACM Technical Symposium on Computer Science Education V*. 1, 603–609.
- [24] Akshay Kolgar Nayak, Yash Prakash, Sampath Jayarathna, Hae-Na Lee, and Vikas Ashok. 2025. Insights in Adaptation: Examining Self-reflection Strategies of Job Seekers with Visual Impairments in India. *Proceedings of the ACM on Human-Computer Interaction* 9, 7, Article CSCW304 (2025), 30 pages. doi:10.1145/3757485
- [25] Akshay Kolgar Nayak, Yash Prakash, Sampath Jayarathna, Hae-Na Lee, and Vikas Ashok. 2026. Contextual Scaffolding and self-efficacy: Supporting computer skill development among blind learners in India. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. 1–21.
- [26] Jonathan Lazar, Aaron Allen, Jason Kleinman, and Chris Malarkey. 2007. What frustrates screen reader users on the web: A study of 100 blind users. *International Journal of human-computer interaction* 22, 3 (2007), 247–269.
- [27] Hae-Na Lee, Sami Uddin, and Vikas Ashok. 2020. TableView: Enabling efficient access to web data records for screen-magnifier users. In *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility*. 1–12.
- [28] Valentyn Melnyk, Vikas Ashok, Valentyn Melnyk, Yury Puzis, Yevgen Borodin, Andrii Soviak, and IV Ramakrishnan. 2015. Look ma, no aria: generic accessible interfaces for web widgets. In *Proceedings of the 12th International Web for All Conference*. 1–4.
- [29] Microsoft. [n. d.]. Seeing AI. <https://www.microsoft.com/en-us/ai/seeing-ai>.
- [30] Peya Mowar, Aaron Steinfeld, and Jeffrey P Bigham. 2026. iTagPDF: Towards finally automating PDF accessibility. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [31] NV Access. 2022. NV Access. <https://www.nvaccess.org/>.
- [32] Cindy Passmore, Alison E. Dobbie, Michael Parchman, and James Tysinger. 2002. Guidelines for Constructing a Survey. *Family Medicine* 34, 4 (2002), 281–286.
- [33] Yi-Hao Peng, Dingzeyu Li, Jeffrey P Bigham, and Amy Pavel. 2025. Morae: Proactively pausing ui agents for user choices. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*. 1–14.
- [34] Yash Prakash, Akshay Kolgar Nayak, Shoaib Mohammed Alyaan, Pathan Aseef Khan, Hae-Na Lee, and Vikas Ashok. 2024. Improving usability of data charts in multimodal documents for low vision users. In *Proceedings of the 26th International Conference on Multimodal Interaction*. 498–507.
- [35] Yash Prakash, Akshay Kolgar Nayak, Nithiya Venkatraman, Sampath Jayarathna, Hae-Na Lee, and Vikas Ashok. 2026. Micro-Behavioral Analysis of Online Shopping Patterns for Blind Users. In *Proceedings of the 2026 Conference on Human Information Interaction and Retrieval*. 295–305.
- [36] Yash Prakash, Akshay Kolgar Nayak, Mohan Sunkara, Sampath Jayarathna, Hae-Na Lee, and Vikas Ashok. 2024. All in one place: Ensuring usable access to online shopping items for blind users. *Proceedings of the ACM on Human-Computer Interaction* 8, EICS (2024), 1–25.
- [37] Yash Prakash, Mohan Sunkara, Hae-Na Lee, Sampath Jayarathna, and Vikas Ashok. 2023. AutoDesc: Facilitating Convenient Perusal of Web Data Items for Blind Users. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*. 32–45.
- [38] Yury Puzis, Yevgen Borodin, Andrii Soviak, Valentyn Melnyk, and IV Ramakrishnan. 2015. Affordable web accessibility: A case for cheaper ARIA. In *Proceedings of the 12th International Web for All Conference*. 1–4.
- [39] Nuzhah Gooda Sahib, Anastasios Tombros, and Tony Stockman. 2012. A comparative analysis of the information-seeking behavior of visually impaired and sighted searchers. *Journal of the American Society for Information Science and Technology* 63, 2 (2012), 377–391.
- [40] Felix Maximilian Schmitt-Koopmann, Elaine May Huang, Hans-Peter Hutter, and Alireza Darvishy. 2025. Towards more accessible scientific pdfs for people with visual impairments: Step-by-step PDF remediation to improve tag accuracy. In *Proceedings of the 2025 CHI conference on human factors in computing systems*. 1–16.
- [41] Joachim Schöpfel, Michael Boock, Behrooz Rasuli, and Brenda van Wyk. 2025. New frontiers of electronic theses and dissertations. *Encyclopedia* 5, 1 (2025), 6.
- [42] Joachim Schöpfel, Eric Kergosien, Stéphane Chaudiron, and Bernard Jacquemin. 2016. Dissertations as data. In *19th International Symposium on Electronic Theses and Dissertations (ETD 2016): "Data and Dissertations"*. sp.
- [43] Yixian Shen, Hang Zhang, Yanxin Shen, Lun Wang, Chuanqi Shi, Shaoshuai Du, and Yiyi Tao. 2025. Altgen: Ai-driven alt text generation for enhancing epub accessibility. In *Proceedings of the 2025 International Conference on Artificial Intelligence and Computational Intelligence*. 78–83.
- [44] Korey J Singleton and Kristine S Neuber. 2020. Examining how students with visual impairments navigate accessible documents. *Journal of Visual Impairment & Blindness* 114, 5 (2020), 393–405.
- [45] Mohan Sunkara, Akshay Kolgar Nayak, Sandeep Kalari, Yash Prakash, Sampath Jayarathna, Hae-Na Lee, and Vikas Ashok. 2025. Adapting Online Customer Reviews for Blind Users: A Case Study of Restaurant Reviews. In *Proceedings of the 22nd International Web for All Conference*. 135–146.
- [46] Mohan Sunkara, Akshay Kolgar Nayak, Sandeep Kalari, Satwik Ram Kodandaram, Sampath Jayarathna, Hae-Na Lee, and Vikas Ashok. 2024. Assessing the Accessibility and Usability of Web Archives for Blind Users. In *International Conference on Theory and Practice of Digital Libraries*. Springer, 203–221.
- [47] Dillip K Swain. 2010. Global adoption of electronic theses and dissertations. (2010).
- [48] Nithiya Venkatraman, Akshay Kolgar Nayak, Suyog Dahal, Yash Prakash, Hae-Na Lee, and Vikas Ashok. 2025. AccessMenu: Enhancing Usability of Online Restaurant Menus for Screen Reader Users. In *Proceedings of the 22nd International Web for All Conference*. 36–47.
- [49] W3C. 2009. WCAG. W3C Web Content Accessibility Guidelines. <http://www.w3.org/TR/WCAG10/>.
- [50] Bryan Wang, Gang Li, and Yang Li. 2023. Enabling conversational interaction with mobile ui using large language models. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–17.
- [51] Wei Xiao and Shuyi Sun. 2020. Dynamic lexical features of PhD theses across disciplines: A text mining approach. *Journal of Quantitative Linguistics* 27, 2 (2020), 114–133.
- [52] Iris Xie, Rakesh Babu, Soohyung Joo, and Paige Fuller. 2015. Using digital libraries non-visually: understanding the help-seeking situations of blind users. *Information Research: An International Electronic Journal* 20, 2 (2015), n2.
- [53] Iris Xie, Rakesh Babu, Hyun Seung Lee, Shengang Wang, and Tae Hee Lee. 2021. Orientation tactics and associated factors in the digital library environment: Comparison between blind and sighted users. *Journal of the Association for Information Science and Technology* 72, 8 (2021), 995–1010.
- [54] Kristin Iyotis. 2008. Electronic theses and dissertation (ETD) repositories: what are they? Where do they come from? How do they work? *OCLC Systems & Services: International digital library perspectives* 24, 2 (2008), 101–115.