Investigating Phishing Threats in the Email Browsing Experience of Visually Impaired Individuals

Emaan Bilal Khan
LUMS
Pakistan

Emaan Atique
LUMS
Pakistan

Mobin Javed
LUMS
Pakistan

ABSTRACT

Phishing, a prevalent cyber threat, persists despite email platform security measures. Visual cues are vital for users to identify phishing, but visually impaired individuals (PVIs) face challenges due to reliance on screen readers. We conduct a qualitative task-based study (n=11) in a Pakistani context, analyzing the interaction of PVIs with phishing emails that target their unique vulnerabilities. Our thematic analysis reveals PVI navigation patterns, highlights challenges in detecting phishing, and provides recommendations for stakeholders.

CCS CONCEPTS

• **Security and privacy → Human and societal aspects of security and privacy**.

ACM Reference Format:
Emaan Bilal Khan, Emaan Atique, and Mobin Javed. 2024. Investigating Phishing Threats in the Email Browsing Experience of Visually Impaired Individuals. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24), May 11–16, 2024, Honolulu, HI, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3613905.3651076

1 INTRODUCTION

With over 300,000 reports during 2022 in the United States alone, phishing has persisted as the world's most common cybercrime [1]. Predominantly, these attacks take the shape of suspicious emails that land in unsuspecting inboxes and often involve the extortion of sensitive information from victims. Email platforms have taken steps to address this challenge, for e.g., the use of security indicators and spam filters. However, numbers indicate that phishing remains an unsolved problem, as the attackers continue to adapt. Studies have in fact reported that over 4.7 million email-based attacks were logged in 2022 alone [2].

A key user defense against phishing is the use of visual cues to determine if an email is legitimate. Downs et al. explore such identifiers in their study, finding that email formatting, out-of-date logos, and inconsistent grammar can play an important role in helping users ascertain the nature of an email [12]. Additionally, confirming the sender's email domain is also a dominant cue used to identify and act against phishing attempts [23].

While these features may be an essential safeguard used by the wider community, for the 2.2 billion globally that suffer from visual impairment [3], the threat of phishing appears heightened. According to the World Heath Organisation, visual impairment, an eye condition that affects the visual system, has four levels of severity: mild, moderate, severe and blindness. Severely visually impaired persons (PVIs), who are also the focus of this study, rely completely on screen reading tools for informed navigation on the web. These tools use speech synthesizers to read aloud visible Internet content that is text-based, a process that is evidenced to be frustrating and difficult, especially if the visited webpage is inaccessible, thereby limiting their defense against online fraud [8].

Related literature in this space, albeit sparse, validates online safety as a real problem for the visually impaired. A study by Kaushik et al. identifies website security assessment to be a challenge for PVIs, and proposes a new extension supporting safer web browsing [16]. Similarly, Yu et al. find that current email browser warnings do not match PVI screen reading habits, and subsequently test an inclusive design to improve warning noticeability [26]. Other work narrows in on the design space, proposing sound-based solutions for defending against security threats [11, 21].

However, there remains a gap in understanding the scope of phishing threats that affect PVI users and their typical interactions with different visual cues such as sender domain, links, and email formatting. This is because studies such as Sonowal et al.'s, which propose a phishing detection model to aid persons with visual impairments lack actual user involvement, both in the design and evaluation process, thus, falling short in incorporating the firsthand experiences and perspectives of actual PVI users [21]. Similarly, while Yu et al.'s work informs insights on existing challenges for PVIs to detect phishing, its exclusive focus on email browser warnings narrows the potential scope of uncovered issues. This leaves room for further work to explore how phishing email content, structure, and other elements beyond email warnings influence PVI behavior.

In addition, importantly, all work in this context is carried out on WEIRD[1] populations, with marked differences in sociocultural, linguistic, and economic backgrounds compared to the non-WEIRD PVI community. This disparity may manifest in the form of distinct security practices, unique email interaction patterns, differences in screen reading tools and their capabilities used by this population, and varying phishing threats driven by local contexts and linguistic differences. Studies have also shown that lower digital literacy significantly increases risk of phishing victimization [14]. These factors ultimately call for research efforts to build a more diverse and representative understanding of the challenges faced by non-WEIRD PVIs in the realm of online email security.

Motivated by these gaps, we aim to study how PVIs interact with diverse email types (varied across email formatting, content,

[1]Acronym for "Western, Educated, Industrialized, Rich and Democratic".

CHI EA '24, May 11–16, 2024, Honolulu, HI, USA

ACM ISBN 979-8-4007-0331-7/24/05
https://doi.org/10.1145/3613905.3651076

attachments and media types), the unique obstacles they encounter in detecting different phishing categories, drawing comparisons between their experiences and their own perceptions of phishing. We particularly focus on sounds-alike phishing threats, which involves manipulation of language to create spoken content that closely resembles legitimate information, an attack tactic disproportionately impacting PVIs. Our research questions are as follows:

- RQ1: How do visually impaired individuals (PVIs) navigate diverse email types using screen readers?
- RQ2: What unique challenges does this group encounter in detecting phishing threats and what gaps exist in their phishing awareness?
- RQ3: In view of PVIs, what improvements can stakeholders, such as email providers and screen reading software companies, make to aid users in better detecting email phishing threats?

In order to answer these questions, we conduct a task-based study with 11 visually impaired email users from Pakistan, where we observe their navigation of a simulated inbox containing phishing emails, analyzing how they choose to process different email types, and what their reading patterns look like. We follow this with a post-task interview to understand their awareness of, experiences with, and defenses against phishing, contextualised with their performance in practice, in order to inform more inclusive platform design.

Overall, our contributions can be summarised as follows:

(1) We conduct the first PVI user study (n=11) focused on phishing email threats in a non-WEIRD context, uncovering unique email navigation and security behavior patterns.
(2) We curate 4 different types of phishing emails, contextualized to our study population, each having different cues that PVIs would struggle to detect. We study the impact of each cue based on user interactions, and identify key threats not systematically explored previously.
(3) Our findings highlight a number of user recommendations for email platform providers and screen reading software companies to address this issue. These reaffirm prior proposed solutions like sounds-based email warnings as well as introduce new ideas such as greater screen reading customization for email browsers.

2 RELATED WORK

Prior work in the space explores the various dimensions of privacy and security concerns surrounding PVIs, identifies current web accessibility gaps, and delves into phishing threats overall, as well as in the context of PVIs. We expand on each of these areas below:

Privacy and Security Concerns of PVIs: There is a wider focus on understanding PVI privacy perceptions, aimed at identifying commonly faced threats and concerns, both physical and online [6, 17, 27]. Through 14 semi-structured interviews, Ahmed et al. [5] identify frequent physical concerns to include eavesdropping of sensitive information, and lack of autonomy while common online considerations surround heightened social media privacy concerns and increased susceptibility to malignant web redirection. Researchers have also studied security and privacy issues for visually impaired people in specific contexts such as Web authentication [13] and cookie notices [10], and a large body of work exists on their interaction with visually assistive technologies [22]. Hayes et al. [15] explore the everyday practices of PVIs to understand the security and privacy implications of the ways people with visual impairments interact with their allies, showing their dependence on external support, and making them wary of who they rely on. Predominantly, research identifies that PVIs harbour heightened concern for their privacy, acknowledging poor design choices that incur increased risk for these users.

Web Accessibility: Current state of work in web accessibility revolves around evaluating websites to assess compliance with policy guidelines like WCAG (Web Content Accessibility Guidelines)[2][20, 24]. Related work also includes devising frameworks that quantify accessibility of these web pages and highlight the design features which contribute to inaccessibility of these web pages [9]. Moreover, studies also shed light on how advertisements can add to various accessibility barriers in the context of people with visual impairments [18]. Findings indicate that advertisements lack accessibility features, which cause disturbance, insecure feelings, as well as user performance decrease, content obstruction, web page backtrack and security issues for persons with visual impairments. Our work builds on the premise that the web proves to be inaccessible for PVIs, narrowing focus on email platforms in order to assess the role their current design may play in increasing risk of phishing.

Protecting People with Visual Impairments from Phishing Emails: While a large body of work focuses on protecting users from phishing emails, prior work focused on PVIs in the context of phishing remains scarce, with room for further research. Specifically focusing on phishing websites, studies indicate that assessing the credibility of a web page is more challenging for those with visual impairments [4]. Another study, focused on browser extensions designed to protect against phishing websites, finds accessibility issues such as color-based security indications, and a lack of shortcut keys, informing the design of Guardlens to help PVIs assess the nature of websites [16]. Blythe et al. [7] focus on phishing emails, interviewing eight individuals with visual impairments, concluding that PVIs are adept at identifying phishing emails due to increased caution, with screen readers aiding in capturing spelling errors. This is a very small sample and is focused on a US population which may influence their phishing awareness. A related study by Yu et al. [26] focuses on email security indicators, assessing the efficacy of current indicators for PVIs, ultimately creating and testing a new indicator design. This work also identifies the gap in screen reading habits and email platform design, but focuses specifically on email warnings. Given that email providers cannot accurately flag all phishing emails and warn users, especially those in multilingual contexts, we explore PVI interaction with various types of phishing emails local to Pakistan. We study how existing results translate to non-WEIRD populations, and compare individuals' perceptions of their phishing vulnerability to their actual susceptibility in practical contexts. Moreover, we particularly focus on sounds-alike phishing in our work, which disproportionately impacts PVIs and remains unaddressed in prior literature from the user perspective.

[2]WCAG is an international standard that outlines how to make web content accessible to people with disabilities, including visual impairments.

3 METHODOLOGY

We use a three step approach to study the threat of phishing for PVIs, shown in Figure 1. We begin with crafting a simulated inbox which contains four different kinds of phishing emails and six benign emails. We curate this inbox on both Gmail and Outlook. In step two, we conduct a task-based user study with 11 PVI participants, recruited via snowball sampling, observing their interaction with the inbox. Finally, we conduct a post-task interview to gauge participants' self reported phishing defense and knowledge.

3.1 Experimental Set-up

3.1.1 User Scenario. Before constructing the phishing emails, we define a user scenario that the emails are premised on. In order to ensure that the target population is familiar with the subject matter, we choose a scenario that is locally relevant. As such, it centers around a fictional person, Imran Hammad, who is the head of the LESCO division of Lahore, a regional electric distribution company. We create the context that Imran has been on leave, so has pending emails in his inbox, which participants will process, posing as Imran. The full scenario can be found in Appendix A.3.

3.1.2 Phishing Email Design. The first step post scenario creation involves the curation of different emails to be utilized as part of the task-based study. We first identify the broad themes common across popular phishing emails. We shortlist the following email types and build each email with associated characteristics that PVI users would find challenging to identify, and are also common to mainstream phishing practices. Images of all emails can be found in the Appendix A.1.

Email 1 (Spearphishing): We choose to curate an email that involves a popular phishing tactic called spearphishing to assess whether PVIs can identify specialized phishing attempts. Spearphishing is a targeted form of phishing attack where cybercriminals tailor their deceptive messages to a specific individual. In this email, we craft a spoofed message, the sender imitating Imran's boss, Rizwan, and asking him to pay an outstanding invoice immediately and discreetly. The phishing cues we incorporate in this email (Figure 3 Appendix A.1) to heighten its suspicious nature include (a) a fake sender email address, which consists of a series of random alphanumeric characters, (b) capitalization of the complete email body, and (c) a sense of unnatural urgency in the language used, calling for discreet actions.

Email 2 (Account Credential Theft): The second email we curate involves the predominant scam category, account theft, whereby individuals are tricked into sharing personal login information for their online accounts. For this, we craft a fictitious Outlook email that suggests unusual account activity, demanding the receiver logs in to a hyperlinked page to prevent account disablement. Cues in this email (Figure 4 Appendix A.1) include (a) a sounds-alike sender email address (doonaughtreply@outlook.com), which will be read as "do not reply" by screen readers, (b) incorrect spelling (i.e., usage of "u/ur" in place of "you/your"), and (c) hyperlink of a crafted phishing website (Figure 7 Appendix A.1), emulating the design of the Microsoft login page, with incorrect grammar and spellings (mikerosoft-account-review.netlify.app). The use of sounds-alike phishing assesses whether it is an additional vulnerability for PVIs, while link insertion evaluates user interaction with spoofed websites in the case that they succumb to phishing emails.

Email 3 (Sensitive Information Extortion): This email employs another prevalent phishing scheme, focusing on sensitive information extortion, where cyber-criminals pose as authoritative bodies to coerce victims into revealing sensitive personal data. In this message, we pose as the Federal Board of Revenue, the national body overseeing taxation. The email alleges income tax discrepancies in Imran Hammad's filing, and urges prompt contact with the Commissioner Inland Revenue to address the issue, including an attached file containing the noted discrepancies, intending to learn more about the victim's personal financial assets. Cues indicating the unusual nature of this email (Figure 5 Appendix A.1) include (a) a spoofed sender email address (fbr.gov.pak@outlook.com), posing as "Fedral Bord of Revenew", (b) highlighted email body, written in colored text, (c) highly pixelated FBR logo, (d) corrupted attached file, and (d) addition of Urdu text.

Email 4 (Malware Installation): We incorporate another prevalent phishing scheme in the next curated email, where the focus is on targeting victims for the installation of malware. In this scheme, cybercriminals often employ deceptive tactics to trick recipients into downloading and installing malicious files onto their systems. Added elements (Figure 6 Appendix A.1) include (a) a spoofed email address (ptcl.technical.help@gmail.com), posing to be the PTCL (a national telecom provider) Support team, (b) unusual font for the email body, (c) inconsistent formatting (indentation, use of bullets), (d) attached Google Drive link containing a zip file labeled "new", and (e) an outdated logo image.

3.1.3 Inbox Creation. To simulate a realistic scenario, six benign emails are crafted, including emails from colleagues and promotional content. One email is image-only to assess the processing of redTen email accounts are registered for each of the crafted emails (both benign and phishing), and the emails are sent to two inboxes (Gmail and Outlook) for the experiment. This is done to ensure coverage of participant email preferences, so participant actions are not influenced by an unfamiliar interface.

3.2 Participant Recruitment

Post pre-testing with one visually impaired person, we initiate online participant outreach, recruiting 11 PVIs, ensuring they regularly use screen readers, and Gmail or Outlook in their daily lives to confirm familiarity with the experimental set up. We further employ snowball sampling to expand the participant pool. The study population includes five individuals aged 18-25, four aged 25-30, and two aged 30-40. Seven identify as completely blind, four as severely visually impaired. Eight prefer Gmail, and three prefer Outlook, with an even split for screen reader usage between NVDA and JAWS, two popular commercial screen reading tools. Participants come from diverse backgrounds, including students, teachers, social media managers, entrepreneurs, and musicians. Significant skew existed in the gender of our participants, with ten who identified as male, and one as female. We acknowledge that future work expanding the sample size to improve this ratio would be beneficial in capturing more differences in our participant demographic.

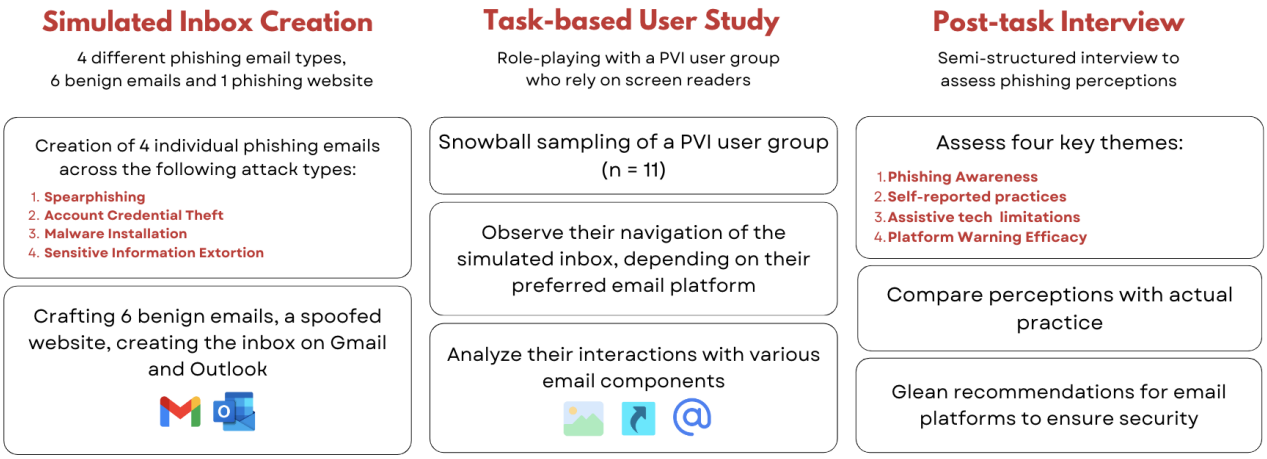

Figure 1: Methodology overview

3.3 Task-based Study

We conducted a task-based user study, virtually on Zoom with 10 participants and in person with 1, ensuring a consistent protocol for both modes to prevent impact on study outcomes. The set-up phase involved sharing mock inbox login credentials based on the participant's platform preference (Gmail or Outlook). Users logged in using their own preferred device (desktop or laptop), screen reader configuration, and browser. Two used Gmail's HTML view, three native email clients, and the rest Google Chrome. Participants screenshared the inbox, and as heads of LESCO, processed pending emails verbally. The study lasted 45-60 minutes for each participant.

Ethics: This study involved a mild degree of deception to prevent raised alertness for phishing, so that we could capture typical PVI security behaviour. As a result, participants were informed pre-recruitment that the study focused on email platform accessibility, and were debriefed about the true purpose immediately after the task-based observation to maintain high ethical standards. Participant consent was sought before the study was recorded and reaffirmed after its completion. All participants were compensated PKR 1000 for their time. Data collected was also anonymized and kept confidential.

3.4 Post-task Interview

Following inbox processing, a post-task interview was conducted to understand users' self-reported phishing habits and perceptions after revealing the study's true objective. This interview covered three main themes: (i) understanding their awareness and prior experience with phishing, (ii) gauging their self-reported security practices, including how they assess email legitimacy and interact with specific content (e.g., images, links), and (iii) assessing their perspective on the current email security landscape as well as expectations from email service providers and screen readers. The interview guide can be found in Appendix A.4.

3.5 Analysis Techniques

To distil key insights from the data, we followed a multistep process. Initially, all interviews were transcribed, and observations from each participant's task-based study and post-task interview were synthesized for analysis. One interview was excluded due to the participant's use of both screen enlargement and screen reading software, which could introduce inconsistencies.

Post-consolidation, we adopted a bottom-up open coding approach, involving two rounds of coding following Saldaña's recommendations [19]. Initially, two researchers independently coded the first five interviews to create a codebook encompassing over 40 lower-level codes, such as "skips email sender address," "ignoring promotional emails," and "not caring about warnings." Meetings were conducted to discuss codes and reconcile any discrepancies. Subsequently, an axial coding analysis was collectively performed to merge similar codes into high-level themes. The resulting themes were "phishing detection," "phishing cues," "email reading patterns," "interactions with non-textual content" (images, links, attachments), "security warning assessment," and "platform recommendations."

4 ANALYSIS AND RESULTS

In this section, we investigate the resulting patterns found post qualitative coding, in order to measure the risk of phishing that people with visual impairments are subject to, specific vulnerabilities for certain phishing cues, and assessment of current security warnings.

4.1 Phishing Detection

We begin with an assessment of the rate and distribution of phishing detection. For this, we define "phishing detection" as the correct identification of the suspicious nature of an email. If an individual chooses to discard an email on the basis of irrelevance (e.g: promotional content) instead of suspicion, we do not regard that as phishing detection.

Figure 2a shows the distribution of phishing detection rate across all participants, confirming that no individual is able to detect all four emails. This indicates that certain phishing scenarios still leave PVIs vulnerable to deceptive tactics. The scenario that is most dangerous is seen to be spearphishing, as seen in Fig. 2b, implying that heightened risk exists for PVIs in the context of targeted phishing attacks. It is important to note that available cues in this email are not noticed, albeit very standard (e.g., sense of grave urgency, financial coercion), given the trust that is placed on the sender name, without confirmation of the origin email. This is

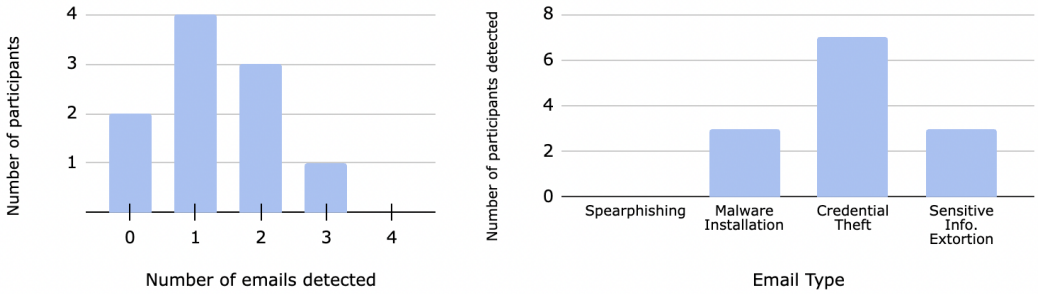

Figure 2: (a) Distribution of phishing detection rate across all participants. (b) Phishing detection rates for different email types.

validated by the following quote: "I can't make the boss angry – I will reply and process the payment immediately" [P5].

We also observe that 80% of users identify at least one phishing email, showing a moderate level of awareness potentially malicious content. This indicates that PVIs are privy to the idea of phishing, which is corroborated by the post-task interview, where all 10 users had knowledge of the concept. However, this awareness does not translate into active and consistent security driven behavior. As such, the steep decline in phishing detection, with only one participant detecting 3/4 phishing emails, and none detecting all, highlights the danger still exists strongly across varied email types, and that PVIs are very likely to fall prey to more nuanced phishing, despite their awareness.

Figure 2b also reveals an important trend. Phishing emails that attempt credential theft have a 70% detection rate among PVIs. Participants raised suspicion immediately, with one participant claiming, "I have gotten such scam emails before" [P10]. Heightened recognition may be attributed to the fact that this phishing category is highly popular, with reports stating 52.8% of phishing attempts in Q1 of 2023 involve credential theft [25]. As a result, there may be widespread awareness of the email type, in turn, resulting in greater user caution. This suggests that more awareness of the kinds of phishing attacks may be an important lever to increase their chance of detection.

Upon conducting Pearson's chi-square test for assessing the role of email type on phishing detection, we find a statistically significant association ($\chi^2(2)$=11.282, p=0.01), indicating the large influence category of email has on the risk that it may pose to PVIs. All in all, we find our results validate that the threat of phishing for PVIs exists, strongly influenced by factors like email type, familiarity with phishing categories, and the level of trust placed in the sender. This stands at odds with their perceived risk of phishing assessed during the post-task interview, where 9 of 10 individuals claim to exercise high caution in the context of phishing attempts. The words of one participant, "I'm extremely careful about these things, that's why I've never been subject to it", incidentally also one of the 2 participants unable to detect any phishing email, mirror the misplaced sense of security PVIs may have when it comes to their phishing defense mechanisms. This underscores the importance of bridging this perception-action gap to better empower PVIs against the evolving landscape of phishing threats.

4.2 Phishing Cues

Now, we discuss participants' perceptions and interactions with the incorporated cues within each email. Our objective is to assess the effectiveness of each element in aiding participants in identifying potential phishing attempts, in order to determine which cues are most helpful in a visually impaired user's context.

4.2.1 Senders' Email Address. Our results reveal that participants often adhere to common email reading habits, neglecting the sender's email address. We find that only 40% of users screen read the sender email address at least once, others instead tending to rely on the sender's name for identification purposes, utilising shortcuts to skip directly to the main email body. This habit is corroborated by Yu et al.'s findings [26], and indicates a potential gap in PVI behavior weakening their phishing defense. Notably, during the spear phishing attempt, all participants exclusively relied on the sender's name, overlooking a spoofed email address containing random letters (i.e. u3u98939209p74993@gmail.com). Visual cues like these may be more noticeable to individuals who can visually inspect on-screen content without relying on screen readers who do not read sender addresses in their typical email reading pattern.

Moreover, we find that only 4 participants demonstrated awareness of registered email domains as an indicator of sender authenticity. These participants raised doubts about the legitimacy of emails associated with unregistered domains, specifically the malware installation and the sensitive information extortion emails.

Exit interviews revealed that participants considering sender addresses in assessing email credibility had prior experiences with phishing attempts, making them more cautious.

4.2.2 Email Format. Next, we measure the role of formatting cues in phishing detection, including grammatical errors, inconsistent font styles, and highlighted text — features often relied upon by users to identify phishing emails.

Upon examination, we find that none were able to identify the format of the spearphishing email. All participants unanimously agreed to process the bank transfer, overlooking the fact that the email was composed entirely in capital letters—a typical indicator of unprofessional communication, unexpected from a boss. When probed during the post-task interview, most individuals agreed that detection of such variations was infeasible using screen readers. A participant commented that certain screen readers like NVDA vary their reading tone for capital text. However, we find that the efficacy of this feature is limited in the typical usage patterns of

PVIs, who are not able to discern this information, given the speed they set for convenient screen reading.

Relying solely on the screen reader's auditory output, participants also failed to recognize misspelled words that sounded similar, such as "donnot" and "doonaught," and "Federal Board of Revenue" with "Fedral Bord of Revenew." Grammatical mistakes like "this is inform you" in the credential theft email was also only noticed by two participants – a cue that may be more conspicuous visually.

Use of slang like "u" and "ur" is again missed by all but one participant. Interestingly, this participant adopted a unique strategy, reading the email character by character, enabling them to identify such errors. However, they acknowledged the impracticality of this approach in typical email reading due to the additional time it entails.

This limitation was also acknowledged across post-task interviews. In the words of one participant, "The screen readers' speech rate is usually very high so it is hard to figure out spelling mistakes, you can't really tell if it's pronounced correctly or not" [P7]. As a result, these cues may be completely missed in the typical usage patterns of PVIs, who have subsequently less information to ascertain the nature of a potentially dangerous email.

4.2.3 Non-textual Content. Next, we assess trends in PVIs' interaction with non-textual content, including links, images, and attachments, typically accompanying the email body and most commonly the malicious feature of phishing emails. Criminals often use hyperlinks to dangerous sites while attempting theft of sensitive data, like account credentials or bank details. Attachments may also contain viruses or malware, intended to target the victim's device.

Links: While analyzing patterns of behavior, we find that users are wary of opening links, which may be attributed to their previous experiences, where three users comment that they or someone they knew had been subject to successful phishing upon opening a link. However, when navigating the emails we find that most users open the Google drive link without waiting for the screen reader to complete its verbal oration, indicating that hyperlinks that are presented as another innocuous link may be opened in haste. One participant commented that it becomes cumbersome to read the full link, which is why they make the decision to open the link before they hear the full content.

Interestingly, when investigating participant behavior with the crafted Microsoft login link in the credential theft email, we notice that two participants copy it to another textbox, and read it character by character in order to determine its validity. This indicates promising and responsible behavior on the part of these users, but the multistep process raises questions about the difficulty this procedure may entail in everyday contexts.

Attachments: Upon analyzing participant behaviors when interacting with email attachments, we find widespread willingness to open attachments without proper verification. 80% of users either indicate they would download the file attached to the malware detection email, or actually download it instinctively while assessing the email's content.

This is kept to be a corrupted file, and worryingly, when a user is not able to open it, they attempt to download it again, assuming error on their part. Even users who adopt a security conscious approach do not feel a hesitation in downloading the attached file, stating that their antivirus software will "ensure security of their device." This trust may be misplaced, and indicates a need for greater user awareness of the limitations of anti-malware software. The interviews also reveal grave threats beyond the crafted emails, whereby two PVIs shared phishing experiences where the extension of files was deliberately renamed to seem non-harmful, such as the renaming of ".exe" files to ".doc.exe", which, in turn, caused malware installation upon download, even resulting in ransomware in one case. These quoted incidents emphasize the urgent need for enhanced user education and awareness regarding file extensions.

Images: Images also remain a commonly inaccessible, and potentially dangerous feature of suspicious emails. Our study findings reveal that accessing this image content, when embedded in emails (like in a benign email we use), or attached, cannot be read or ascertained without the use of OCR, which requires that users download and scan the image. Inadvertently, PVIs that want to learn more about an image have to download unverified material on their devices, another potential threat to their safety. In particular, one participant recognized his dependence on those around him to describe image content, saying "I often rely on my friends to tell me what it [the image] is about" [P3], raising important concerns about resulting compromise of PVI privacy. This also speaks to the limitations of screen readers, withholding the ability from users to grasp or access all available information in emails, given the dominant occurrence of images.

4.3 Security Warning Assessment

While none of the emails we sent triggered any browser's security indicators, we explore the perceived efficacy of these safeguards in conversations with users. Interestingly, we find that three users mention never having encountered any such warnings in their typical email browsing experiences, whereas four point out their limited efficacy due to the likelihood of false positives, where genuine emails are often mislabeled as spam. The remaining participants mentioned their value in redirecting spam content, but did not comment on their role in the case of phishing emails.

This observation underscores a significant challenge in relying solely on browser security indicators, as users may develop skepticism or overlook potential threats when they become accustomed to false positives, diminishing the effectiveness of these warning systems. It also reveals a gap in their current design and implementation, given we do not find them on our phishing emails across both email providers.

4.4 Participant Recommendations

Several recommendations emerged from participants' experiences and perspectives to enable better phishing detection.

Screen Reader Improvements Users suggested improvements in screen reader capabilities, including the ability to differentiate between fonts and styles, as well as more efficient detection of word capitalisation to enable more informed email navigation. Introducing customization options for adapting the order of screen reader output (i.e., read sender, then email, then body, instead of the fixed template arrangement of the web content which skips to email body without reading sender address) was also suggested.

These improvements have not been uncovered in prior research, and underscore the importance of involving users in designing and reviewing accessibility tools to ensure they meet inclusion standards. In the same vein, while some screen readers may already include similar customizations, the lack of PVI awareness about these features points to the need for better education and training initiatives to maximize the utility of existing accessibility options.

Email Browser Enhancements Recommendations for email providers included the need to flag incorrect spelling and grammar in received emails, similar to the way unrecognised words are highlighted while emails are drafted. Users also advocated for warnings to be presented as pop-ups or read aloud rather than banners, ensuring immediate attention. However, in the words of a participant, "sound based warnings, if frequent, will be annoying" [P2], calling for more variety in proposed solutions, driven by user recommendations. Additionally, suggestions to filter emails based on sender addresses and clearer verbal indication of attachment types, especially for those deceptively labelled with a bogus extensions, like "example.png.exe" were important recommendations to improve phishing detection.

Participants also presented insightful suggestions to improve the accessibility of the email browsing experience. These included the recommendation to automate alternative descriptions of attached or embedded images within emails, citing similar practices by platforms like Meta, which would help avoid the practice of downloading images to read them through external OCRs. Additionally, participants suggested decluttering web elements in the accessible version of the email browser to improve overall usability. This would enable better detection of phishing cues and allow for a more personalized user experience based on individual preferences. This is increasingly important as the two participants who use the HTML version of Gmail due to its accessibility point out that Gmail is discontinuing it permanently in January 2024, indicating the urgent need for tailored email browser designs to ensure PVI convenience and safety online.

Taken together, these recommendations provide a rich roadmap for technology practitioners to develop more robust and user-centered email security solutions for visually impaired users. The urgent need for accessible email browser designs in light of Gmail's HTML discontinuation also underscores the importance of continuous user research and the value of participatory design with PVI communities. Ultimately, these user driven suggestions can serve as a solid foundation for both academia and industry to build on as they work to ensure safer and more empowering online experiences.

5 STUDY LIMITATIONS

Despite the valuable insights gained from this study, there are certain limitations that should be acknowledged. Firstly, the sample size, while diverse in terms of age, occupation, and visual impairment severity, was relatively small and skewed by gender. A larger and more diverse participant pool would enhance the generalizability of the findings. Additionally, the study primarily focused on participants who were familiar with screen readers, Gmail and Outlook. While this specificity ensured a targeted investigation, it limits the extrapolation of results to users employing different email platforms, browsers, or assistive technologies. Furthermore, the use of a simulated set-up might not fully capture the real-world dynamics of handling emails, as participants were aware that the study involved assessing email accessibility. This awareness might have influenced their behavior, potentially leading to a more cautious approach than their typical email interactions. Despite these limitations, the study lays a foundation for future research in understanding the challenges faced by visually impaired individuals in detecting phishing threats through email.

6 CONCLUSION AND FUTURE WORK

In summary, as the first non-WEIRD PVI user study on email phishing, we explored their interactions with diverse email types, the unique obstacles they encounter in detecting different phishing categories, and highlighted gaps between phishing awareness and active detection. Spearphishing emerged as a significant vulnerability, emphasizing the need for greater PVI awareness. Participants exhibited differing levels of scrutiny when assessing phishing cues, revealing deficiencies in evaluating sender email addresses and formatting, particularly for PVI populations dependent on screen readers. Additionally, interactions with non-textual content like images exposed potential risks given their need to be downloaded before they could be assessed. There also appeared widespread skepticism of the efficacy of browser indicators due to false positives. User-driven recommendations, including customizable screen reading order and accessible browser options, in turn, offer actionable insights for stakeholders to enhance security for PVIs.

All in all, the results of this study set the stage for further exploration, with an expanded participant sample to validate the generalizability of findings. Moreover, user recommendations point to the possibility of co-designing interventions to help PVIs detect phishing threats with greater accuracy, for more informed and safer email navigation.

REFERENCES

[1] [n. d.]. https://www.ic3.gov/Media/PDF/AnnualReport/2022_IC3Report.pdf

[2] [n. d.]. https://apwg.org/trendsreports/.

[3] [n. d.]. https://www.who.int/publications/m/item/increasing-eye-care-interventions-to-address-vision-impairment

[4] Ali Abdolrahmani and Ravi Kuber. 2016. Should I Trust It When I Cannot See It?: Credibility Assessment for Blind Web Users. https://doi.org/10.1145/2982142.2982173

[5] Tousif Ahmed, Roberto Hoyle, Kay Connelly, David Crandall, and Apu Kapadia. 2015. Privacy Concerns and Behaviors of People with Visual Impairments. 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, 3523–3532. https://doi.org/10.1145/2702123.2702334

[6] Tousif Ahmed, Patrick Shaffer, Kay Connelly, David Crandall, and Apu Kapadia. 2016. Addressing Physical Safety, Security, and Privacy for People with Visual Impairments. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, Denver, CO, 341–354. https://www.usenix.org/conference/soups2016/technical-sessions/presentation/ahmed

[7] Mark Blythe, Helen Petrie, and John A. Clark. 2011. F for fake: four studies on how we fall for phish. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.

[8] Yevgen Borodin, Jeffrey P. Bigham, Glenn Dausch, and I. V. 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)*. Association for Computing Machinery, 1–10. https://doi.org/10.1145/1805986.1806005

[9] John Breton and Abdelrahman Abdou. 2023. Applying Accessibility Metrics to Measure the Threat Landscape for Users with Disabilities. In *Networks and Distributed System Security Workshop on Measurements, Attacks, and Defenses for the Web (MADWeb) 2023*.

[10] James M. Clarke et al. 2023. Invisible, Unreadable, and Inaudible Cookie Notices: An Evaluation of Cookie Notices for Users with Visual Impairments. *ArXiv abs/2308.11643* (2023).

[11] P. Datta, A. S. Namin, K. S. Jones, and et al. 2021. Warning users about cyber threats through sounds. *SN Appl. Sci.* 3 (2021), 714. https://doi.org/10.1007/s42452-021-04703-4

[12] Julie S. Downs, Mandy B. Holbrook, and Lorrie Faith Cranor. 2006. Decision strategies and susceptibility to phishing. *Proceedings of the second symposium on Usable privacy and security - SOUPS '06* (2006). https://doi.org/10.1145/1143120.1143131

[13] Ahmet Erinola, Annalina Buckmann, Jennifer Friedauer, Aslı Yardım, and M. Angela Sasse. 2023. "As Usual, I Needed Assistance of a Seeing Person": Experiences and Challenges of People with Disabilities and Authentication Methods. In *2023 IEEE European Symposium on Security and Privacy Workshops*.

[14] Roderick Graham and Ruth Triplett. 2016. Capable guardians in the digital environment: The role of digital literacy in reducing phishing victimization. *Deviant Behavior* 38, 12 (Nov 2016), 1371–1382. https://doi.org/10.1080/01639625.2016.1254980

[15] Jordan Hayes, Smirity Kaushik, Charlotte Emily Price, and Yang Wang. 2019. Cooperative privacy and security: learning from people with visual impairments and their allies. In *Proceedings of the Fifteenth USENIX Conference on Usable Privacy and Security (SOUPS'19)*. 1–20.

[16] Smirity Kaushik, Natã M. Barbosa, Yaman Yu, Tanusree Sharma, Zachary Kilhoffer, JooYoung Seo, Sauvik Das, and Yang Wang. 2023. GuardLens: supporting safer online browsing for people with visual impairments. In *Proceedings of the Nineteenth USENIX Conference on Usable Privacy and Security (SOUPS '23)*. 361–380.

[17] Daniela Napoli. 2018. *Accessible and usable security: Exploring visually impaired users' online security and privacy strategies*. Ph. D. Dissertation. Carleton University.

[18] Ab Shaqoor Nengroo and K S Kuppusamy. 2019. 'Advertisements or adversetisements?'—An accessibility barrier for persons with visual impairments. *Comput. J.* 62, 6 (2019), 855–868. https://doi.org/10.1093/comjnl/bxy104

[19] Johnny Saldaña. 2015. *The Coding Manual for Qualitative Researchers*. Sage.

[20] Karen Schnell and Kaushik Roy. 2021. Website Privacy Notification for the Visually Impaired. In *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*.

[21] Gunikhan Sonowal. 2020. A Model for Detecting Sounds-alike Phishing Email Contents for Persons with Visual Impairments. In *2020 Sixth International Conference on e-Learning (econf)*. 17–21. https://doi.org/10.1109/econf51404.2020.9385451

[22] Abigale Stangl, Emma Sadjo, Pardis Emami-Naeini, Yang Wang, Danna Gurari, and Leah Findlater. 2023. "Dump it, Destroy it, Send it to Data Heaven": Blind People's Expectations for Visual Privacy in Visual Assistance Technologies. In *Proceedings of the 20th International Web for All Conference*.

[23] Bradley W. Weaver, Adam M. Braly, and David M. Lane. 2021. Training users to identify phishing emails. *Journal of Educational Computing Research* 59, 6 (2021), 1169–1183. https://doi.org/10.1177/0735633121992516

[24] Kathrin Wille et al. 2016. Measuring the Accessibility Based on Web Content Accessibility Guidelines. In *2016 Joint Conference of the International Workshop on Software Measurement and the International Conference on Software Process and Product Measurement (IWSM-MENSURA)*. 164–169.

[25] John Wilson. [n. d.]. Phishing trends and tactics: Q1 of 2023. https://www.tripwire.com/state-of-security/phishing-trends-and-tactics-q1-2023

[26] Y. Yu, S. Ashok, S. Kaushik, Y. Wang, and G. Wang. 2023. Design and Evaluation of Inclusive Email Security Indicators for People with Visual Impairments. In *2023 IEEE Symposium on Security and Privacy (SP)*. 2885–2902. https://doi.org/10.1109/SP46215.2023.10179407

[27] Yuhang Zhao, Yaxing Yao, Jiaru Fu, and Nihan Zhou. 2023. {"If} sighted people know, I should be able to {know:"} Privacy Perceptions of Bystanders with Visual Impairments around Camera-based Technology. In *32nd USENIX Security Symposium (USENIX Security 23)*. 4661–4678.

A APPENDIX

A.1 Phishing Emails

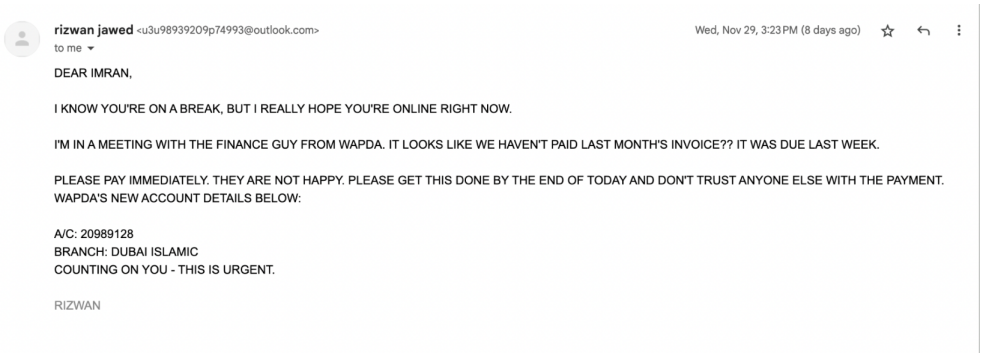

Figure 3: Spearphishing (Email 1)

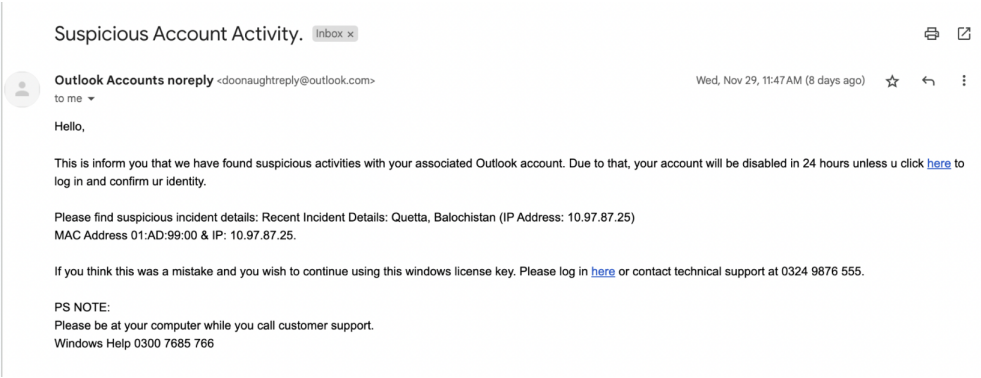

Figure 4: Account Credential Theft (Email 2)

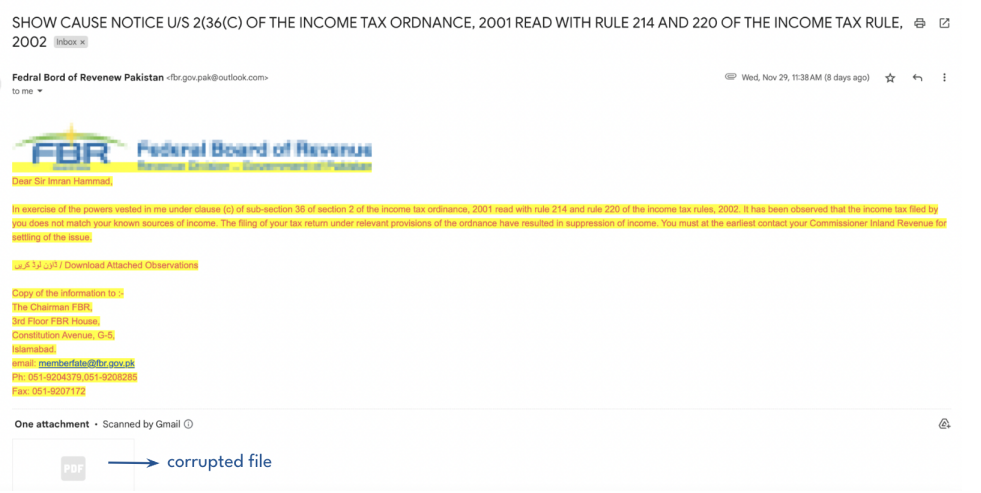

Figure 5: Sensitive Information Extortion (Email 3)

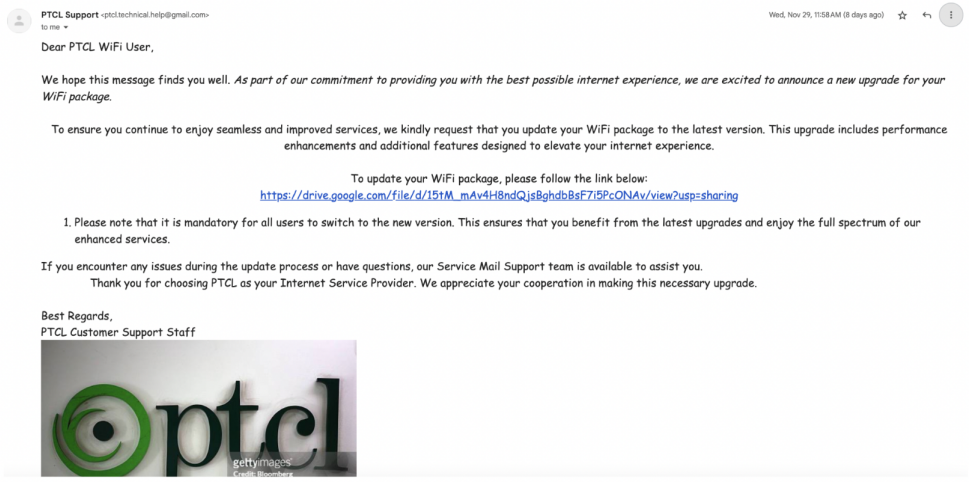

Figure 6: Malware Installation (Email 4)

A.2 Phishing Website

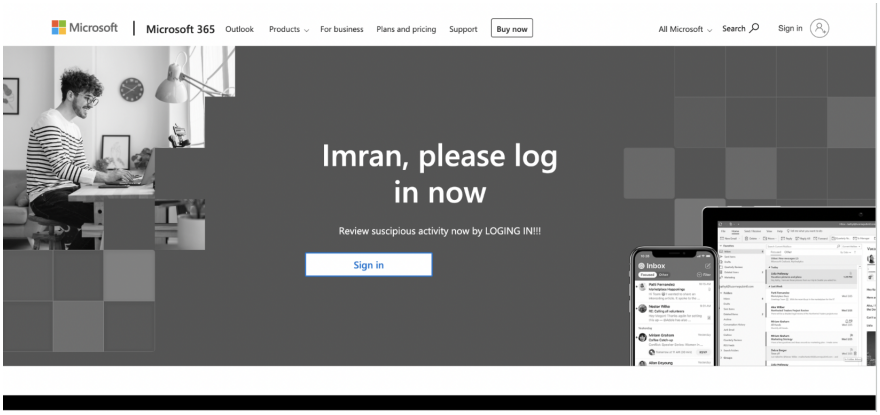

Figure 7: Sample phishing website used in credential theft email (Email 2)

A.3 Task-based Study Scenario

"Your name is Imran Hammad, and you are the head of the LESCO division of Lahore. Your boss is Rizwan Javed, who you directly report to. Emaan Bilal is your technical lead, and Hassan Ahmed, from the LESCO IT team, also directly reports to you.

Today is Monday, November 27, 2023. You were on vacation last week and did not check your emails. You have already arrived at the office at 9:00 a.m. to complete the tasks at hand. You have half an hour to process the emails you have received, as you will no longer be in the office from 9:30 a.m. onwards, because you have a number of other appointments. You will not be able to reach any of your co-workers or business partners, as the official start of work is not until 10:00 a.m. The order in which you process them is entirely up to you. Your task is to choose what you think is the best applicable action for each email. Please make sure to process all the emails in the inbox!"

A.4 Post-task Interview Guide

Thank you for all your enthusiasm! We have a few closing questions to make better sense of the risk of phishing to PVIs particularly.

Segment 1: General Awareness and Knowledge

Basic Understanding

- What does the term "phishing" mean to you?

Personal Experience

- Have you or anyone you know ever been a victim of a phishing attack? If yes, could you briefly describe the incident and its impact?

Segment 2: Practices

Email Recognition

- How do you determine if an email is legitimate or potentially a phishing attempt?
- What are some red flags you look for in emails?

Links and Attachments

- How cautious are you when clicking on links or opening attachments in emails?
- Are there any specific types of links or attachments that make you more suspicious?

Segment 3: Accessibility

- Do Gmail/Outlook security warnings help in detecting emails that are unreliable?
- What do you suggest in terms of email providers' and screen reader responsibilities for a more secure email browsing experience?