

WebAlly: Making Visual Task-based CAPTCHAs Transferable for People with Visual Impairments

Zhuohao Zhang

University of Illinois, Urbana-Champaign

Haolin Yuan

John Hopkins University

Sauvik Das

Georgia Institute of Technology

Zhilin Zhang

University of Illinois, Urbana-Champaign

Natã Barbosa

University of Illinois, Urbana-Champaign

Yang Wang

University of Illinois, Urbana-Champaign

Abstract

Task-based visual CAPTCHAs are a significant accessibility hurdle for people with visual impairments (PVI). What if PVIs could transfer task-based visual CAPTCHAs to a helper to solve? How might PVIs want such a system configured in terms of from whom they would solicit help and how they would compensate this help? To answer these questions, we implemented and evaluated a proof-of-concept assistive transfer system — WEBALLY — that makes task-based CAPTCHAs transferable by allowing PVIs to source just-in-time, remote control help from a trusted contact. In an exploratory, role-play study with 10 pairs of participants — a PVI and a friend or a family member — we asked participants to use WEBALLY in four different configurations that varied in source of help (friend vs. stranger) and compensation (paid vs. volunteer). We found that PVIs liked having WEBALLY as an additional option for solving visual CAPTCHAs, when other options that preserve their independence fail. In addition, many PVIs and their friends felt that using the system would bring their relationship closer. We discuss design implications for transferable CAPTCHAs and assistive transfer systems more broadly, e.g., the importance of complementing rather than replacing PVIs’ existing workflows.

1 Introduction

Motivation. Large swathes of the web remain inaccessible for the 285 million people with visual impairments (PVI) [41]. For instance, CAPTCHAs (Completely Automated Public Turing tests to tell Computers and Humans Apart) are commonly used to authenticate users in numerous day-to-day web

surfing tasks [33] (e.g., registering new accounts, leaving comments on social media, and completing financial transactions), yet were rated as the most problematic item for PVIs in a global study by WebAIM [51].

CAPTCHAs are inaccessible for PVIs because they require users to engage in complex visual-processing tasks. For example, today, task-based visual CAPTCHAs, such as Google reCAPTCHA [46] and GeeTest [21], are widely used across the web and require users to perform high-precision operations such as selecting a subset of images from the gallery, or dragging a slider to solve a puzzle. These CAPTCHAs are challenging if not impossible for PVIs to solve independently.

Existing solutions. Prior work has explored a number of solutions to make CAPTCHA-solving easier and more accessible. Most commonly, PVIs use audio CAPTCHAs instead. However, prior work has found that audio CAPTCHAs are disproportionately hard for PVIs relative to how hard visual CAPTCHAs are for people without visual impairments — they are significantly slower and require more attention and memory-capacity [15]. Other solutions to help PVIs in the short-term include automated CAPTCHA solving services (e.g., WebVisum [56]), but these solutions work only for simple visual CAPTCHAs in which people are asked to identify distorted letters and numbers. There are also CAPTCHA solvers (e.g., Anti-CAPTCHA [7]), but these pose security risks — they require users to install software with dangerous system-level permissions. In short, while existing solutions do help PVIs solve or bypass CAPTCHAs in some cases, there are still many other cases in which they fail PVIs.

To help PVIs in overcoming day-to-day web accessibility hurdles outside of CAPTCHAs, crowdsourcing and friend-sourcing methods have shown great promise. Many PVIs use remote assistance services or ask friends around them to directly help [58]. While these methods require interdependence, crowdsourcing and friendsourcing can help PVIs with these accessibility challenges that they might encounter in today’s web. Prior art, such as BeMyEyes [10] and VizWiz [16], have explored connecting PVIs in need of help with remote assistance from sighted helpers to, for example, answer ques-

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

USENIX Symposium on Usable Privacy and Security (SOUPS) 2021.
August 8–10, 2021, Virtual Conference.

tions about surroundings or request for vocal-guidance on using inaccessible interfaces. However, these applications are limited to descriptive guidance, which limit their utility in solving task-based CAPTCHAs.

There are also trade-offs to sourcing help from the crowd or from friends. While crowdsourcing might pose privacy and security risks for remote control — in which a crowd helper (stranger) takes over the PVI’s system to solve CAPTCHAs on their behalf — friendsourcing better aligns with existing workflows for PVIs: PVIs often seek assistance from their friends and family members to overcome accessibility challenges in the physical world [1]. However, friendsourcing can also lead to privacy problems as users might not want to expose their ongoing activities to friends. Friendsourcing can also reduce monetary cost of using paid crowd workers. However, friendsourcing can be slower and less reliable than crowdsourcing [9], and PVIs might want to avoid burdening their social connections with frequent requests for help [47].

Research questions. In short, PVIs commonly encounter task-based visual CAPTCHAs that frustrate and encumber their use of the web, yet existing solutions fall short of supporting their needs. With an overarching goal of designing inclusive privacy/security tools [55], we posit that an *assistive transfer system* that allows PVIs to solicit just-in-time help where a willing helper directly solves an outstanding reCAPTCHA challenge on the PVIs’ behalf may be beneficial, if seen as a “last resort” when other options to retain independence have been exhausted. However, the design space of such assistive transfer systems have not yet been explored, and there are a number of open questions regarding how PVIs might use such a tool and how it might be configured — e.g., whether and in what contexts the help should be sourced from friends or strangers and whether they would prefer the help to be free or paid. To bridge this gap, we designed and implemented WEBALLY, a proof-of-concept assistive transfer system that allows PVIs to transfer the solving of task-based CAPTCHAs to others, and used WEBALLY as a design probe in an exploratory user study to answer these research questions:

- **RQ1:** What are both PVIs’ and helpers’ general impressions towards using assistive transfer systems like WEBALLY?
- **RQ2:** What are the perceived privacy and security risks for PVIs when transferring CAPTCHAs to others?
- **RQ3:** What factors influence PVIs’ preferences in configuring from whom to source help?
- **RQ4:** What factors influence PVIs’ preferences towards compensating helpers?
- **RQ5:** What is the perceived impact, of using assistive transfer systems like WebAlly, on the social relationship between PVIs and friends from whom they solicit help?

User study. We conducted a within-subjects, two-by-two (crowdsourcing vs. friendsourcing, paid vs. free) lab study (over Zoom video conference) with 18 participants (10 PVIs and 8 sighted friends) to answer our research questions. We recruited participants in pairs (one PVI and one friend who served as a remote helper), and had PVIs use WebAlly to request their helper to solve Google reCAPTCHAs for them. To simulate the different configurations of our study, we had participants role-play — a technique commonly employed in usable security research [48]. For example, while each PVI participant had the same helper (their friend) in all conditions, we had participants envision themselves in a situation where they would need to solicit help from the crowd/stranger or a friend, and where they would need to pay for this help or not. Our data included notes taken by researchers as participants engaged in the study tasks, and post-study exit interviews. Given the limitations of our study, we note that our key contribution is *less* about WEBALLY and its evaluation in and of itself, but *more* about — the knowledge of designing transferable CAPTCHAs and assistive transfer systems more broadly — gained through the design and evaluation of WEBALLY with stakeholders, as is common in HCI design research [63].

Findings. We found that while transferring CAPTCHAs requires interdependence, both PVIs and helpers appreciated having a system as a last-resort alternative to other accessibility solutions (e.g., audio CAPTCHAs). PVIs felt that WEBALLY could mitigate the privacy and security risks entailed by transferring a task to helpers by intelligently cropping to only task-relevant parts of the PVIs’ screen, but some still had concerns about sourcing help from others altogether. Helpers, on the other hand, had some concerns over requests from strangers in the form of an open link. We also uncovered four important factors that may affect PVIs’ perceptions towards using an assistive transfer system for CAPTCHAs: the type of webpage that embeds CAPTCHA (e.g., whether it contains private browsing data), the use case (e.g., whether it is financial-related), helper availability, and impact on requester-helper relationships. In addition, we also explored PVI users’ personal preferences towards compensating the helpers.

Contributions. Our work has two main contributions: (1) We introduced and explored the design space of assistive transfer systems for task-based CAPTCHAs, and (2) We implemented a proof-of-concept assistive transfer system, WEBALLY, and conducted an exploratory evaluation with both PVIs and helpers to synthesize design insights for assistive transfer systems in the context of Google reCAPTCHAs.

2 Related Work

2.1 PVI with CAPTCHAs

CAPTCHAs are designed to distinguish humans from robots. They deter hackers from abusing online services and are served millions of times a day [30]. Traditional CAPTCHAs

usually pose a visual challenge, like recognizing images, words, or numbers out of specific images. These interactive tasks are meant to be simple for human users. However, CAPTCHAs are notoriously inaccessible for folks with visual, physical, cognitive, or auditory disabilities [14, 26, 37, 49]. For PVIIs who use screen readers, specifically, visual CAPTCHAs pose a hurdle that often cannot be overcome without relying on other external help — be it a friend or a service designed to bypass CAPTCHAs.

As an alternative, audio CAPTCHA is more accessible for PVIIs. However, audio CAPTCHAs are not always available on many websites, and current audio CAPTCHA designs have proven to be difficult and time-consuming for PVIIs throughout several research studies [15, 32, 34, 35]. To improve security against speech recognition algorithms [17], the audio file provided to users are usually speakers saying words at randomly spaced intervals with background noise. These interferences challenges both automated agents and human users [32, 52]. Many existing research studies have also tried to increase the accessibility of audio CAPTCHAs. Fanelle et al. designed four novel audio CAPTCHAs to increase accuracy and speed [24]. Jain et al. proposed *reCAPGen*, a system that uses automatic speech recognition for generating more usable and secure audio CAPTCHAs [32]. They all explored how users (especially PVIIs) can independently solve audio CAPTCHAs. Prior work also provided many examples of directly breaking CAPTCHAs. Some early research leveraged image and pattern recognition techniques to break visual CAPTCHAs [20, 38, 60]. More recent research also provided various types of hackings towards task-based CAPTCHAs like Google reCAPTCHA [8, 36, 50, 61]. There are also many paid CAPTCHA solving services like Anti-CAPTCHA [7] and Buster [25]. Although these techniques could be easily adopted in browser extensions or system-level applications to hack CAPTCHAs directly for PVIIs, the original purposes of these research are still aimed for improving the CAPTCHA’s security by revealing how they can be hacked. Additionally, using hacking services like Anti-CAPTCHA would introduce privacy and security issues. Users will need to download browser extension files directly from their website rather than installing from official stores, and users are required to edit their computer’s registry to make the tool work.

2.2 Privacy and Security Concerns of PVIIs

As online resources have been more and more available and accessible for users with visual impairments, there is a trend towards empowering PVIIs to protect private information and their online security. Gurari et al. introduced the first visual privacy dataset originated from PVIIs, revealing a challenge of understanding and protecting their privacy needs [27]. The dataset also includes information that can be easily captured on PVIIs’ computer screens. As crowdsourcing remote assistance services like BeMyEyes and Eyecoming [42] have been

widely used by PVIIs and make their lives easier, researchers have also investigated how these services would raise privacy and security risks [3, 5, 6, 59]. Akter et al. conducted a study to understand privacy concerns when PVIIs use camera-based assistive technologies [6]. Ahmed et al. took another angle and studied the information sharing preferences of sighted bystanders of assistive devices [3]. Existing research has also shown that PVIIs have strong security and privacy concerns in using CAPTCHA [2, 4, 22, 28, 31]. Holman et al. identified their top 10 security challenges and CAPTCHA has been listed as the top one challenge [31], which poses a challenge of how to help PVIIs solve these small tasks like CAPTCHA without compromising their privacy and security.

2.3 Sourcing help for PVIIs

Socio-technical researchers conducted many studies on collaborative systems. Traditional crowdsourcing has proved a convenient way to get answers quickly from the crowd. The VizWiz smartphone application allows visually-impaired users to send visual questions to sighted crowd workers and get answers soon [16]. However, such services can be limited due to the cost of the paid crowd workers, which might add extra and unexpected burden to PVIIs [19]. Friendsourcing could also help users solicit answers and assists from friends via online social network services, and the answers are often from more trustworthy and tailored to their interests than using a search engine [40]. Traditional online social network sites used for these include Facebook and Twitter [13, 39, 40, 47]. For PVIIs specifically, AbdraboTarek et al. proposed an assistive tool for blind users to friendsource help for daily activities via smartphone and Twitter [1]. Brady et al. studied PVIIs’ perceptions of social microvolunteering via Facebook answering visual questions on behalf of blind users [19].

Crowdsourcing and friendsourcing have their unique advantages and disadvantages in many aspects. For example, differences exist about compensation and response rates and potential impact on social relationships. Zhu et al. studied the effects of extrinsic rewards and monetary payments to further investigate how friendsourcing would impact PVIIs’ social relationship with their friends [62]. Other research also revealed how these rewards might undermine the original motivation that drives friendsourcing activity and change the perceived relationship between people [29, 43, 53, 54]. In addition, independence is often considered as a goal in assistive technologies [11]. Even sometimes the goal is not explicitly stated, the researchers agree that “all accessible computing approaches share a common goal of improving independence, access, and quality of life for people with disabilities” [57]. However, as Bennett et al. pointed out, interdependence is also valuable because the interactions between people with disabilities and their allies are often two-way and mutually beneficial [11]. In our work, we aimed to use a novel collaborative method as a probe to explore these different design spaces. We also looked at both PVI user side and helper’s side,

which has not yet been fully investigated by other researchers.

3 Design Considerations

We began by identifying the challenges PVI's experience with task-based CAPTCHAs — we enumerated existing solutions to help PVI's overcome these CAPTCHAs, investigated how existing solutions fall short of PVI's' needs, and uncovered how this transferable method could play a role in helping PVI's solve these tasks. Then, we synthesized several design goals to support PVI's with an accessible tool to solve Google reCAPTCHA, one of the most commonly used task-based CAPTCHA on the web.

3.1 Design Challenges

C1 - Providing PVI's with more direct manipulation Existing tools such as BeMyEyes and Eyecoming provide PVI's with remote assistance services: e.g., providing descriptive guidance on how to operate an interface and navigation guidance via smart glasses. However, these collaborative assistance services are limited to providing indirect help; PVI's rely on helpers' textual or verbal guidance, either synchronously or asynchronously, to solve the task on their own. While this type of assistance is helpful, fosters independence, and is widely used by PVI's, it can become challenging when the task requires precise hand-eye coordination (e.g., moving the mouse and clicking specific areas). One opportunity to address this challenge is to afford a remote helper direct control of the PVI's system to solve the task-based CAPTCHA on behalf of the PVI. However, it is still challenging to make remote control assistance secure and accessible for PVI's.

C2 - Protecting privacy and security while helping While remote control assistance could help PVI's overcome task-based CAPTCHAs, it might also bring privacy and security issues: as many PVI's may be unable to receive visual feedback or otherwise monitor helpers' behaviors, remote helpers could perform malicious actions on PVI's' devices without their awareness. The helper could also become aware of what the PVI is trying to do online, or be able to see sensitive personal information that may be present on the PVI's screen. Thus, there is a need to protect PVI's' privacy and security as they receive help through remote assistance systems and services — both for indirect descriptive guidance (when PVI's usually need to point their cameras to the computer screen) and direct remote control.

C3 - CAPTCHA restrictions In CAPTCHA design, there is usually a trade-off between security and user experience. In achieving its original purpose of differentiating humans from bots, task-based CAPTCHA can be challenging even for humans as a result of making it more robust against bots. For example, Google reCAPTCHA, the most common type of task-based CAPTCHA, has a solving time limit of two minutes. It also expires within one minute before submission,

such that users must complete and submit a form protected by reCAPTCHA before it expires, lest they have to solve another reCAPTCHA challenge. Remote assistance services usually take more than two minutes to post requests, find volunteers, synchronize with the helper, get help, and get notifications when the session is complete. Often the case is that PVI's need to wait for someone to answer their requests. It is naturally challenging to source help in a short amount of time before the current CAPTCHA expires.

3.2 Design Goals

To address these challenges, we highlighted several goals that we identified as essential for designing an efficient and accessible CAPTCHA-solving tool for PVI's. Our high-level design goals were to integrate social support, reduce human effort, source help efficiently, protect PVI's' privacy, and still maintain the security utility of CAPTCHAs — differentiating between humans and bots.

G1 - Limited remote control To provide PVI's with more direct help and maintain their privacy and security at the same time, our goal is to design a limited, sandboxed remote control system in which helpers are restricted in the actions they can perform and the screen information they can see. Specifically, helpers should only be able to perform actions necessary to solve the CAPTCHA, and should only be able to see parts of the PVI's screen that is relevant to the CAPTCHA. In this case, helpers cannot access any sensitive information or perform other actions on the PVI's' personal devices.

G2 - Simplicity of use Task-based CAPTCHAs like Google reCAPTCHA have time limits, after which they expire and a new visual challenge is issued. One typical assistive system often includes: (1) send the PVI's' request, (2) synchronize with a helper, (3) have the helper complete the task. For a fast and simple solving experience, the user interface for the helpers should be simple and straightforward so they can minimize the completion time to avoid expiration.

G3 - Accessible in usage To ensure that our remote assistance tool is accessible, our goal is to make all its functions available via keyboard shortcuts. Another design goal to enhance accessibility is to provide audio feedback at every stage of helping, to notify PVI's about the current state of the task and what the helper is doing.

3.3 System Design

Guided by design goals, we implemented WEBALLY — a proof-of-concept system that connects PVI's with their friends when they encounter Google reCAPTCHAs. WEBALLY creates an interactive screenshot of the reCAPTCHA and sends it to the helper. This interactive screenshot serves as a canvas on which helpers' can perform actions to solve the reCAPTCHA (e.g., clicking on tiles, dragging UI elements); these actions, in turn, are reflected on the PVI's screen. However, access to the

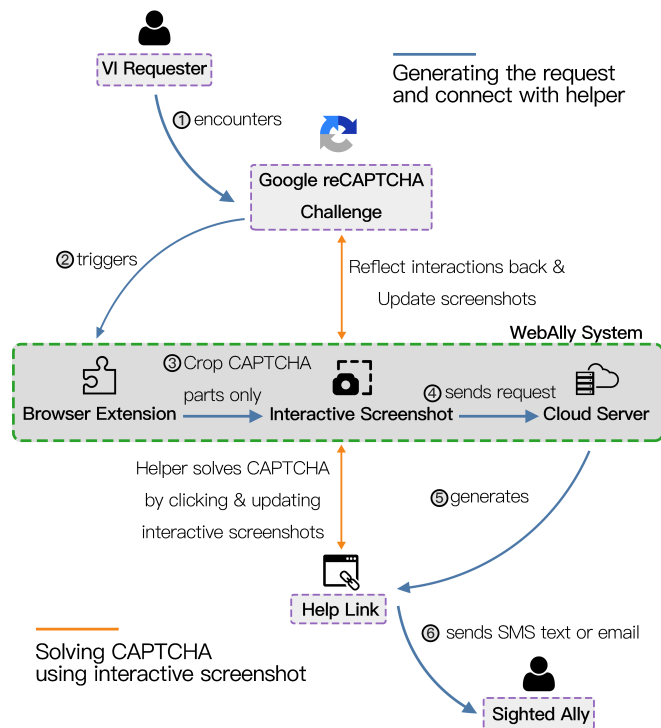


Figure 1: Workflow of WebAlly: PVI triggers the tool and sends the request, and the tool builds a channel between the PVI and the helper via an interactive screenshot: Reflect helper’s clicking on requester’s screen and update screenshots back to helper’s interface until finished.

source device is totally reconstructed — helpers see only task-relevant screen information, and only pre-specified interactions (e.g., clicks) will be reflected and simulated on the PVI’s screen. In the following sections, we will use the terminology “requester” to refer to users with visual impairments sending requests for help, and “helper” to refer to their friends and family offering help. The WEBALLY system stores cropped screenshots and helpers’ contact information only temporarily — i.e., for the duration of the transfer task.

3.3.1 System Overview

WEBALLY is implemented as a browser extension, written in JavaScript and executable on Chromium-based browsers (such as Google Chrome, Microsoft Edge, Opera and Brave). We also incorporated OpenCV to pre-process images, and WebSocket as a channel to transmit messages in real time. The workflow contains a one-way request from the requester and a synchronous collaboration process between the requester and the helper (see Figure 1). We have open-sourced the source code for WEBALLY ¹.

¹<https://gitlab.engr.illinois.edu/salt-lab/webally>

3.3.2 Interface Details

The requester interface is simplified to be accessible for PVIs. First, the requester will enter and store the contact information for a helper using a keyboard shortcut to activate the function. The requester can then activate the extension and send the request to the preset helper using an editable keyboard shortcut. The WEBALLY system then takes a screenshot of the current browser tab and uses the template-matching feature in OpenCV.js to crop the screenshot down to just the region that contains the Google reCAPTCHA task. Only the cropped image will be sent to the helper. If the reCAPTCHA task is successfully solved, the requester can end the collaborative session and continue the task at-hand, e.g., submitting a form. If the helper failed to pass the test in time, the requester could ask again via the same keyboard shortcut.

The helper will receive SMS texts or emails with a URL to a secure page. The helper can then complete the requested task on their own browser. The helper’s interface contains instructions on what to do and the cropped interactive screenshot. To reflect the helper’s actions on their screen into the requester’s screen, WEBALLY record, transmit, and simulate the helper’s clicks on the interactive screenshot via WebSocket. As the reCAPTCHA interface updates the tile images after each click, the system also detects the updated part, crop it using OpenCV, and reflect these changes back on the interactive screenshot presented to helpers. Thus, the helper’s experience mimics how they might solve a CAPTCHA for themselves.

4 User Study

While WEBALLY is fully functional and can be used in practice, it is not a finalized product — it is a design probe that we presented to PVIs in order to model their perceptions of and configuration preferences for transferring reCAPTCHA challenges to remote helpers. To that end, we conducted a two-by-two within-subjects study with 18 participants (10 PVIs and 8 helpers) to evaluate WEBALLY and answer our research questions. Our study was IRB-approved. We had PVIs and their helpers roleplay using WEBALLY to transfer a Google reCAPTCHA task in four different scenarios that varied in source-of-help (stranger or friend) and compensation strategy (paid or voluntary). While we report on some descriptive quantitative findings, we note that our main findings are qualitative — our goal was less to comparatively evaluate different conditions, and more to understand PVIs perceptions of WEBALLY under different configurations.

4.1 Participants

We recruited participants in pairs: one PVI and one helper who the PVI considered a close friend or family member. We recruited 18 participants (see Table 1), including 10 participants (6 males, 4 females) with visual impairments (referred to as requesters R1-R10) and 8 participants (3 males, 5 females)

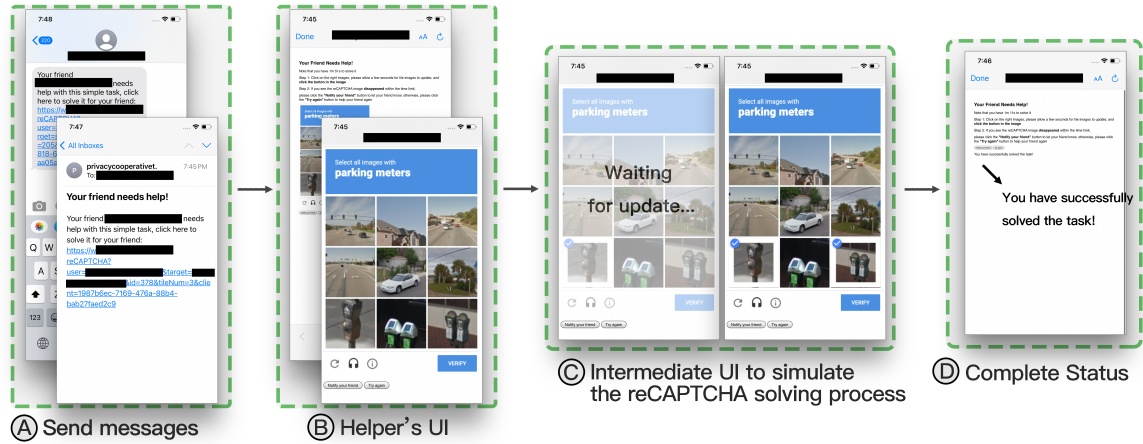


Figure 2: User Interface from the helper’s side: Requester sends message, helper opens the link, solve the CAPTCHA as usual, and click “VERIFY” to complete (The solving process might take multiple visual challenges to pass).

Table 1: Participant demographics, including requesters’ age range, gender identity, self-described visual ability, screen reader of use, and their relationship with the helpers.

PID	Age	Gender	Visual ability	Screenreader	Helper	Gender	Relation
R01	25-34	Male	Blind	NVDA	H01	Male	Brother
R02	18-24	Male	Blind	NVDA	H02	Female	Sister
R03	55-64	Female	Deafblind	NVDA	H03	Male	Brother
R04	25-34	Female	Blind	JAWS	H04	Female	Brother
R05	25-34	Male	Blind	JAWS	H05	Female	Friend
R06	25-34	Male	Low Vision	NVDA	H06	Female	Friend
R07	25-34	Male	Low Vision	NVDA, JAWS	H07	Male	Friend
R08	35-44	Male	Blind	JAWS	H08	Female	Friend
R09	45-54	Female	Blind	JAWS	H08	Female	Friend
R10	35-44	Female	Blind	JAWS	H08	Female	Friend

without visual impairments (referred to as helpers H1-H8 just for simplicity). Among PVI, one self-identified as deafblind, two as low-vision, and seven as blind. One helper (H3) self-identified as having some hearing impairment. For the screen readers they were using, one was using both NVDA [45] and JAWS [44], four were using NVDA, and five were using JAWS. For the study sessions, requesters R1-R7 were in pairs with helpers H1-H7 accordingly in our study, and H8 was a mutual friend to requester R8-R10, who joined the study with them for three times. For their relationship, one helper is the sister, three helpers are the brothers, and four helpers were friends of the corresponding PVI.

4.2 Apparatus

We used the WEBALLY prototype to conduct the study. We asked requesters to install WEBALLY on their Chrome browsers before the study. Instead of using a real website for testing WEBALLY, we used a demo website (<https://www.google.com/recaptcha/api2/demo>) which contains a login form simulating what users would encounter in real settings. The reason we are not using real websites is that,

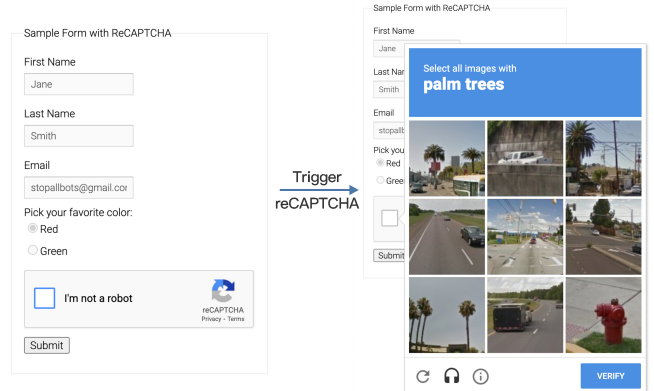


Figure 3: The demo website to trigger Google reCAPTCHA.

unlike real websites, the demo website is designed to *stably trigger* the Google reCAPTCHA task to ensure PVI users would face this particular task in the study.

For helpers, we asked them to have a device (either computer or mobile devices) through which they could join the meeting and receive SMS messages and/or emails.

4.3 Procedure and Data Collection

We conducted a lab study to explore the design space of WEBALLY. While a field study could be helpful in evaluating near-final design concepts in ecologically valid contexts, a lab study allowed us to capture rich, qualitative data on users’ acceptance of and perceived feelings towards WEBALLY under different experimental configurations in an exploratory and controlled way. We do not see WEBALLY as a near-final design concept. Rather, it is a design probe and sensitizing

concept that we implemented to evaluate with stakeholders in order to synthesize design knowledge, as is common in HCI design research [63]. As such, we elected to run a lab study.

The participant sessions were all conducted remotely via video conference calls. Before the sessions, we confirmed that participants had regular access to a laptop or desktop computers, and that they used Mozilla Firefox or Google Chrome to surf the web in their daily life. All sessions lasted about an hour, including a post-study interview to learn the requesters' and the helpers' feedback separately.

At the beginning of each study session, two researchers and two participants (the requester and the helper) would join the online meeting. Participants were introduced to the purpose of the study and the study procedures. In some cases, the requesters did not or failed to install the tool in advance. The researchers then helped them install the tool via the online meeting and guided the requesters to type the helper's information into the WEBALLY's interface.

Scene Setup and Role Explanation

After the preparation and installation, the researchers would divide the two participants into two breakout rooms to simulate remote collaboration (i.e., they did not need to be physically co-located to use the system). The two researchers who helped conduct the study went into the each of the two breakout rooms in order to observe helper and requester behaviors and answer their questions. Splitting requesters and helpers up into two separate breakout rooms helped approximate a real-world scenario in which a PVI would need to solve a Google reCAPTCHA task and choose to source remote help from friends. After the researchers and the participants settled in different rooms, the researcher in the requester room (referred to as Researcher 1) introduced the tasks and asked the participants questions from a pre-study questionnaire (A.1) about their experience with CAPTCHAs and Google reCAPTCHA. The questionnaire was designed to understand how participants generally solve reCAPTCHAs and the challenges faced in solving these task-based CAPTCHAs.

After asking participants about their prior experiences with CAPTCHAs, researchers explained the scene and role setup. In the study, we asked helpers to play the role of both friend and stranger (i.e., crowd worker). This ruse was made more believable by the fact that requesters and helpers were separated from one another during the study in order to simulate the remote collaborative setting. As such, our participants did not necessarily know the exact identity of who might have been helping them or requesting their help for a given reCAPTCHA task. PVI requesters were told that they will transfer the request to either their friend or an unknown crowd worker in different scenarios. Helpers were told that they will receive a request from either their PVI friend or a stranger who is also using our tool. Thus, while we employed role play, participants had reason to believe their roles were, in fact, true

— strengthening ecological validity, though still a limitation.

Similarly, to strengthen ecological validity from the helpers' perspective, helpers were instructed that the requests from their PVI friends or a stranger may come at any time, and that they could do whatever they pleased in the meanwhile rather than waiting for WEBALLY requests. When their assistance was requested, they would be notified via SMS or email — just as they would in real settings when they are not necessarily prepared to help their friends exclusively.

After explaining scene and role setups, the lab study began. Researcher 1 asked the requester to imagine that they are under one of four different scenarios. We had a 2 x 2 within-subjects experimental design with two factors: Source of help (ally vs. stranger), and compensation for help (free vs. paid). In the helper's room, the researcher (referred to as Researcher 2) also introduced the different configurations to the helper and asked them to behave accordingly under different scenarios (e.g., imagine that a blind person whom you do not know asks your for free, voluntary help to solve a reCAPTCHA).

In the study, the researchers randomized the order in which the four configurations were presented to reduce order effects. Under each configuration, researchers asked participants to use WEBALLY to solve a reCAPTCHA task together. Broadly, the PVI asks their helper for assistance, the helper solves the reCAPTCHA on his/her own screen, which would be automatically transmitted back to the helper through the limited remote control functionality we implemented in WEBALLY. To reliably trigger a reCAPTCHA challenge under each configuration, requesters used the aforementioned demo website. Then the requester would send the challenge to the helper through a keyboard shortcut. The helper could receive the help request message via SMS or email (see Figure 2 A). The message would contain a broad description of the requested remote assistance task (see Figure 2 D) and a web link containing the interactive screenshot of the Google reCAPTCHA (see Figure 2 B, C). In ideal scenarios, each configuration contained just one task-solving process. However, if the helper failed to complete the task within the time limit and the reCAPTCHA expired, the researchers asked participants to repeat the process again under the same configuration.

After completing all of the four configurations, the researchers conducted an exit interview asking about participants' detailed experiences with the system and their preferences among the four configurations. The questions include general feedback, suggestions on improvement, which configuration they chose as their favorite and why, their privacy and security concerns in details, and how the request might alter the relationship dynamics between requesters and helpers.

4.4 Analysis

Upon participant consent, we video-recorded all the online sessions and took notes from the procedure and the interview.

We transcribed the videos and used thematic analysis [18] to qualitatively analyze the study. Broadly, our analysis was driven by our core research questions and covered perspectives of both PVI and helpers. For PVI, our codes cover prior experience with solving CAPTCHAs (e.g., tools used or methods for sourcing help), their general feelings about WEBALLY, concerns on using WEBALLY and its potential privacy/security risks, their nuanced context-based preferences with respect to crowd/friendsourcing and paid/unpaid versions of WEBALLY, their desire to use WEBALLY in the future (willing to install and recommend), and their perception of how WEBALLY might affect their relationship with friends and family. For helpers, our analysis covered their overall feeling and suggestions, willingness and general availability to help PVI or even a broader user group with/without compensation, their choice on different configurations to help, and related feelings.

5 Results

5.1 RQ1: General Impressions and System Performance

After participants used WEBALLY under all study conditions, we asked about their general impressions of using the system.

All participants reported that they liked the tool but also identified its pros and cons. For example, many (9 out of 10 requesters and 7 out of 8 helpers) thought the solving process was effective, and they were excited about the overall idea of transferring complex CAPTCHAs to others and how the tool was easy and effective. For instance, R1 praised WEBALLY as a “really bright idea.” R2 said “really appreciate[d] this invention” because “many websites do not have audio CAPTCHAs, and many audio ones just do not work.” Many helper participants also thought positively about the tool. For instance, H6 commended that “it is simple and helpful, and I feel great to help someone.” H4 cited privacy benefits: “it is great especially from a privacy perspective” in reference to the privacy-preserving image cropping features.

However, some users (3 out of 10 requesters and 4 out of 8 helpers) also pointed out limitations to task transference. Most importantly, two PVI were concerned that their friends and family would not be available to solve a CAPTCHA before it expires in two minutes. For instance, they hesitated in sending request messages to their friends when the timing would be inconvenient — e.g., late at night. This result, while unsurprising, is a fundamental limitation to making tasks transferrable to friends and other social connections; transference, thus, must be considered a complement to existing accessibility solutions — a “last-resort” option that can be relied upon when all other options to retain independence have been exhausted. R8 also noted that “it is limited to Google CAPTCHA, and it takes some time. It also expires sometimes.” Indeed, the overhead implicit in task transference may often

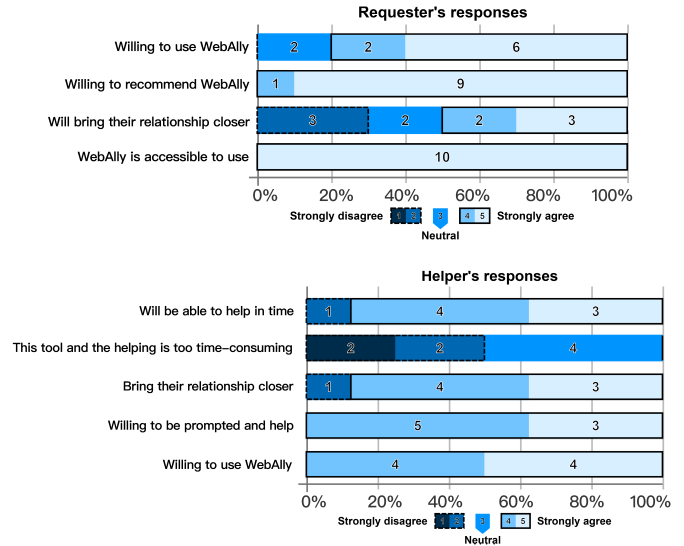


Figure 4: Requesters’ and helpers’ responses on the acceptance of the tool, possible change in their relationship, and the tool’s accessibility level.

be a hurdle in transferring tasks with short fuses — especially those that are dynamic and require round-trips (e.g., reCAPTCHA tasks in which tiles get updated). However, while WEBALLY was designed for Google reCAPTCHAs, it is easy to generalize the core concept of task transference to other accessibility challenges as well.

Among sighted helpers, some (3 out of 8) also expressed concerns about overhead and delay. H1 mentioned that “the server response time could be improved for a bit” and H2 said “sometimes I need to wait for a while for images to update and I got confused when I clicked other squares too quick.” While network latency is unavoidable, WEBALLY could be improved to more gracefully handle this latency.

Participants also provided insights about WEBALLY relative to existing tools. Prior to this study, most of our PVI participants used remote assistance services or CAPTCHA-solving tools. They mentioned the similarities and differences between WebAlly and tools like BeMyEyes and WebVisum. The results showed the uniqueness of WEBALLY, which could adopt helpers’ direct visual perception and corresponding action into solving task-based CAPTCHAs. Services like BeMyEyes provided a crowdsourcing solution to gather descriptive guidance. WebVisum is one of the earliest methods designed for solving visual CAPTCHAs, but many participants mentioned that it stopped maintenance long ago, and they have been relying on friends or family members physically around them to help them solve Google reCAPTCHAs.

Similar to prior work, most requesters (9 out of 10) mentioned that audio CAPTCHAs are not always accessible, and

they will still need help from a person. Based on our pre-study interview, a typical flow for PVIs to solve a Google reCAPTCHA was: (1) try to continue with their task without solving the CAPTCHA (e.g., if they were already logged into their Google accounts); (2) try to solve the audio version; and (3) if failed, ask nearby friends to solve the visual CAPTCHA for them. Half of the requesters mentioned that they could not solve the challenge if there were no friends around. This result suggests the potential value of assistive transfer systems.

We also recorded the time duration and success rate of solving CAPTCHAs as complementary data. However, it is important to note that since the helpers were present and ready to solve transferred CAPTCHAs in our lab study, the time to solve CAPTCHAs through WebAly might be longer in practice. In our lab study, the WEBALLY tool performed well under the four different configurations. We consider a CAPTCHA task to be successful if a PVI user could pass a Google reCAPTCHA by having a helper solve the CAPTCHA via the interactive screenshot within the time limit. Overall, for all the first-time trials in all 10 study sessions, the success rate was 60%, with most failures being due to the 2-minute time restriction imposed by Google reCAPTCHA. On the first trial, helpers were not yet familiar with the tool, resulting in delays. However, as helpers gained familiarity in later trials, their speed and success improved. Indeed, the overall success rate across all task sessions was 88% (SD=0.10). Moreover, all helpers were able to solve a transferred CAPTCHA challenge the second time around if they failed the first time (likely due to their being immediately available and prepared).

Most times, the helpers could successfully solve the task within one minute. The average solving time for each task (if the task was successfully solved) was 37.9 seconds (SD=12.65). For the three study sessions (8th-10th) with the same helper (requesters R8-R10 and helper H8), the solving time decreased (27.25 and 31.5 seconds) and the success rate increased (100% for both), suggesting that trained helpers can complete the transferred tasks quickly and reliably.

5.2 RQ2: Privacy and Security Perceptions

To understand users' privacy and security perceptions of WEBALLY, we asked our participants open-ended questions about concerns that arose in their use of WEBALLY under different configurations, e.g., sourcing help from friends versus strangers.

Privacy and Security Perceptions of The Tool At first, requesters were not aware of the image cropping feature of WEBALLY because the cropped interactive screenshots were only available to the helpers. Without this knowledge, 3 out of 10 requesters (R3, R4, R8) proactively asked researchers if the helpers (irrespective of whether they were friends or strangers) could see their entire screen as would be the case with traditional remote control assistance. After learning that

WEBALLY crops what helpers can see to only task-relevant portions of the UI (e.g., the reCAPTCHA challenge), all requesters found this feature useful in protecting their privacy and security.

Privacy and Security Perceptions of Transferring Tasks

Participants also varied in their perceptions of transferring online tasks more generally. Four out of ten requesters (R4, R8, R9, R10) proactively mentioned privacy and security considerations, but also stated that privacy and security were secondary concerns [23] — they cared more about “*getting the job done*” (R4). Other requesters expressed concerns about privacy and security-related risks when asking strangers on the Internet for help, especially if the task context — i.e., the action that the CAPTCHA was authenticating — was a monetary transaction or signing up for a new account. Although they felt that volunteers on the Internet might have good intentions to help them, some PVIs still had some concerns — R8, for example, said “*you can never be too careful.*”

Requesters' privacy and security concerns varied based on whether they were sourcing help from friends or strangers. When sourcing help from strangers, requesters were most concerned about privacy and security in financial task contexts. They also expressed concerns about what volunteer strangers could access on their computers. When sourcing help from friends, requesters had more concerns about whether their helper could see private information such as browsing history. We asked requesters about concerns they had about using friendsourcing to solve online tasks more generally, where cropping exact section out of the complex screen contents may not always be successful. They expressed a common preference that they do not want their friends to see the whole screenshot and know which website they are visiting (R8: “*I don't want my friends to know where I am looking, but I don't care if strangers see it. They don't know me anyway.*”). In contrast, requesters were unconcerned if crowd workers could access their browsing histories. In short, help source appears to affect requesters' privacy/security concerns for assistive transfer systems.

Helpers' privacy and security perceptions of the tool

From the helpers' perspective, when we asked them to imagine they were helping PVIs that they did not know personally, some helpers (H3, H6) mentioned that they would not trust some SMS messages containing a link from an app on which they were not pre-registered — as they might be phishing links. They suggested that the tool should also provide a channel for helpers to register beforehand (like BeMyEyes) even though they do not need to install the extension to help PVIs. A registration process would give helpers some confidence and trust in the tool when receiving messages.

5.3 RQ3: Factors Affecting Preference on Source of Help

We found that participants' preferences for sourcing help from close social connections (friendsourcing) or strangers (crowdsourcing) varied based on four factors:

- **Factor 1: Page content:** What content is on the webpage in which the CAPTCHA is embedded? Is there sensitive information?
- **Factor 2: Use case:** What is the PVI doing? What task is being authenticated by the CAPTCHA? Is the action sensitive or security-related?
- **Factor 3: Helper availability:** Is it a good time to ask for help?
- **Factor 4: Impact on social relationship:** Does the PVI feel like they are burdening their helpers?

Most participants (7 out of 10 requesters) preferred sourcing help from their friends/family for small favors such as solving CAPTCHAs. They mentioned that asking friends for help would make them feel more secure than asking strangers, and that requesting small favors such as solving a CAPTCHA would not bother friends much (e.g., R2: “*Yes I would rather ask friends since it’s very simple for them,*” Factor 4). We further discuss how assistive transfer systems might impact social relationships in Section 5.5. Participants also expressed a preference for friendsourcing when the use-case was security sensitive, such as a financial transaction (e.g., R4: “*I feel a little uncomfortable when sending this [a CAPTCHA embedded in a money transfer use case] to strangers.*” Factor 2). The three other requesters (R7, R8, R10) who preferred crowdsourcing expressed concerns about bothering their friends (e.g., R7: “*I don’t want to interrupt my friend when it’s midnight*” Factor 3) and privacy issues (e.g., R8: “*I don’t want my friends to know where I am looking*” Factor 1) as discussed in the last subsection. We also found that participants would like to have the option of choosing between friend- and crowdsourcing on a case-by-case basis.

5.4 RQ4: Compensation Preferences

Three out of seven requesters who preferred friendsourcing mentioned saving money as a key rationale for their preference — CAPTCHA tasks are small and simple enough for friends. In comparison, four out of seven requesters still preferred to compensate their friends for helping them, even for small favors such as helping with solving a CAPTCHA.

We also found that nearly every requester (9 out of 10), even the ones who preferred to receive help for free, would prefer to compensate their helper with non-monetary rewards such as “*a cup of coffee.*” They believed paying their friends a small amount of money would make them “*feel weird,*” or that it

would be an “*insult.*” R2 mentioned that a subscription service would also be acceptable for both requesters and helpers since a routine and fixed payment would cause less embarrassment between requesters and helpers.

For the three out of 10 requesters who preferred crowdsourcing, all preferred to use a paid service rather than a free tool. Some requesters (R8, R9) were already using paid remote assistance services to help them with any technical issues they encounter while using computers. These requesters expressed that they would trust the crowd workers more when they paid for the service as the helpers would be “*trained or professional workers.*” (R8)

From the helpers' standpoint, most preferred to be compensated with a non-monetary award. For solving a small task like Google reCAPTCHA, they would also settle for completely voluntary work with no payments or rewards. We also asked whether helpers would be willing to help with strangers' requests — most mentioned that they would be willing, at least in theory. Only one helper (H3) mentioned that helping strangers might be overwhelming because “*you need to pick up random messages.*” Most helpers mentioned that they would not feel bothered by a small number of PVIs with a small number of requests, e.g., 3-5 friends requesting fewer than 5 times per day. In terms of compensation, most helpers would not expect getting paid much or getting paid at all because “*solving a CAPTCHA only takes seconds*” (H8).

5.5 RQ5: Impact on Social Relationships

Surprisingly, most participants (5 out 10 requesters and 7 out of 8 helpers) thought that an assistive transfer system like WEBALLY would bring requester and helper closer in their relationship (see Figure 4). H3 believed this tool could give them “*more contact opportunities,*” and there is “*value of knowing that I helped my sister.*” Some participants thought the tool would not change their relationship since they were already friends, and there were not “*many additional interactions*” (R5). Also, R5 pointed out WEBALLY's social value in raising sighted people's awareness that PVIs need help in these everyday tasks.

We also asked questions about request boundaries and limits for helpers. Most helpers indicated a number between 5 to 8 times a day that they would feel comfortable solving tasks for friends or strangers who need help. If incoming requests exceeded this number, helpers expressed that they would feel bothered; if these requests came from a friend, it would potentially negatively impact their relationship. In a pre-study interview, we asked PVIs how many times they needed help solving CAPTCHAs in their weekly web use — the reported number was less than 10 times weekly, well within the helpers' reported boundaries. While these numbers are speculative and would need to be validated in a field deployment, these results suggest that friendsourcing may be a viable option for assistive transfer systems that help PVIs solicit just-in-time help for reCAPTCHA tasks.

6 Design Implications

Based on our study findings, we synthesized a number of design implications for designing transferable CAPTCHAs and assistive transfer systems, more broadly.

6.1 CAPTCHAs

Consider Making Inaccessible Tasks Transferable Ideally, CAPTCHAs would simply be more accessible, eliminating the need for assistive transfer systems like WEBALLY altogether. However, given the known and longstanding accessibility issues with CAPTCHAs and their alternatives, we suspect that sweeping changes to improve the accessibility of security challenges will be slow in coming. In the interim, CAPTCHA designers might consider making it easier for assistive transfer systems to work. For example, WEBALLY would likely be too slow unless a friend or helper was on standby. CAPTCHA designers might slightly increase the time allowed for a CAPTCHA to be solved if a transfer request is initiated, for example.

The Interplay between Security and Accessibility CAPTCHAs were originally designed to distinguish humans from online bots. The specific ways of making CAPTCHAs more accessible for PVIIs might have security implications — for example, if bots could pose as PVIIs in order to trigger these transfer requests. This issue is more prominent in crowdsourcing than in friendsourcing contexts — presumably, friends would need to be pre-registered and only accept requests from those they personally know. In a crowdsourcing context, fees associated with the service may discourage bots from utilizing such requests. For voluntary crowdsourcing services, the onus should be on the service that facilitates such task transference in doing due diligence (e.g., only registered users can send requests with a time-based limit).

6.2 Assistive Transfer Systems

Our paper focused on exploring the design space of assistive transfer systems — i.e., a system that assists PVIIs by soliciting just-in-time help from friends or helpers — for Google reCAPTCHA. However, our findings offer broader implications for the design of assistive transfer systems.

Complementing, not replacing, existing workflows Given a choice between a CAPTCHA challenge they could solve themselves and a perfect assistive transfer system, PVIIs would likely choose the former. Thus, assistive technologies that support independent use of computing devices should be the ultimate goal. However, the modern web is a far cry from being fully accessible for PVIIs. At least in the short-term, our research suggests that there is value in allowing accessibility hurdles to be transferred to pre-registered friends and

crowdworkers so that PVIIs can have a “last resort” option when they have exhausted options to overcome the hurdle independently. While still interdependent, the use of online assistive transfer systems like WEBALLY could remove the need for the helper to be physically present — often how PVIIs obtain help from trusted allies — for everyday challenges. We note, however, that interdependence is not inherently bad and could also open new design possibilities for assistive technologies that empower PVIIs [11].

Reducing Latency with a Helper List and Speculative Recruitment Most of the PVIIs who participated in our study reported having more than one close friend who was available to help. An active friend list would help increase the chance of tasks getting picked up and solved within the time limit. An assistive transfer system should distribute load across many willing friends and/or helpers, and iterate through the list if there is a delay in response. One could also imagine the use of speculative and/or proactive recruiting, à la Bernstein et al.’s Crowds in Two Seconds approach [12]. When a PVII navigates to a website in which a CAPTCHA request is likely to occur, an assistive transfer system might pre-emptively request a friend or helper to be on standby. Of course, care will need to be taken to ensure the PVIIs’ privacy preferences are respected — they should be offered an informed choice. Another possible option is that helpers can indicate their availability status (e.g., “free,” “busy”). Perhaps assistive transfer workflows could also be integrated with social networking websites and instant messaging apps, eliminating the need for a separate setup and effectively using PVIIs’ existing social connections.

Compensation preferences Our findings suggest that PVIIs and their friends, alike, prefer non-monetary compensation for assistive transfer systems like WEBALLY, mirroring Zhu et al.’s findings in prior work [62]. Thus, in practice, designers should consider creative alternatives to payment, e.g., a small gift as a token of appreciation. In crowdsourcing contexts, PVIIs seemed to prefer a paid subscription service over transactional micro-payments.

New Opportunities for Social Interactions Contrary to our initial expectations, some of our participants mentioned that occasionally sourcing help from friends to overcome accessibility hurdles would be a good excuse to catch up with that friend. Moreover, PVIIs’ friends echoed this sentiment, stating that they felt good in the knowledge that they helped their friend. While a field study is necessary to validate this effect in practice, this result suggests that assistive transfer systems have the potential to help maintain or even improve social relationships between PVIIs and the friends from whom they source help. Such systems might offer new opportunities for meaningful social interactions.

7 Limitations and Future Work

First, our design probe was only limited to the Google reCAPTCHA challenge and was implemented as a Chrome browser extension. The present research did not explicitly cover other CAPTCHAs, transferable tasks or browsers. However, we believe that the Google reCAPTCHA is one of the most popular CAPTCHAs, which present a common challenge for PVI users. Our tool implementation could also be extended to explore other CAPTCHAs and browsers.

Second, we only asked the participant pairs to solve CAPTCHAs using our tool rather than testing additional ways to solve Google reCAPTCHAs. This was mainly because (1) our transferable approach was meant to complement rather than replace existing solutions, (2) our study goal was to explore the transferable task design space, and (3) we felt that asking PVI participants to directly solve CAPTCHAs in addition to using our tool could be exhausting for them. We also avoided direct comparisons with crowdsourcing assistive tools like BeMyEyes because the afforded functionality is significantly different as solving CAPTCHAs requires more than descriptive guidance.

Third, to explore different design configurations of transferable CAPTCHAs, we asked PVI requesters and sighted helpers to imagine that they are in hypothetical scenarios where (1) the helper is an ally (family members or friends) or a stranger/crowd worker, and (2) the helper gets paid or not for solving the CAPTCHA. However, this role-play still has limitations in ecological validity because for instance, it did not capture ally’s availability (e.g., in practice, they might not be available at the time when PVI users need help). A field trial can better represent the realities but it is less suitable for our study goal of exploring the design space of transferable tasks rather than testing the effectiveness of a final system. We also considered having a PVI participant’s ally participant serve as another PVI participant’s stranger/crowd worker. However, we were not able to recruit/schedule two pairs of participants to do the study at the same time due to people’s different schedules. In fact, scheduling a pair of PVI participant and the ally participant to conduct the study together was already very challenging.

Another limitation worth mentioning is that Google reCAPTCHAs will be passed if users are already logged in their Google accounts. However, if the users choose to remain private (e.g., use incognito mode), they will face CAPTCHAs much more frequently. We recognize that PVIs might be in a position where they need to make extra effort to maintain their privacy, and we provide WEBALLY as an additional option when they face such inaccessible challenges.

WEBALLY’s current implementation may open the door to a few security risks. For example, since requests are sent to helpers as URLs via SMS, one can imagine a new vector for phishing unsuspecting helpers. Malicious requesters might also use WebAlly to circumvent CAPTCHAs for free. A field-

ready implementation could mitigate these risks by requiring requester registration and authentication, allowing helpers to whitelist from whom they can receive requests, and imposing daily request quotas to avoid abusive use.

Finally, the assistive transfer system approach requires interdependence between PVIs and their allies, which might affect PVIs’ perceived independence in doing daily tasks. It is important to note, however, that even without WEBALLY, PVIs often ask for help to bypass visual CAPTCHAs. An alternative direction is designing new mechanisms that replace visual CAPTCHAs, for instance, better audio CAPTCHAs that do not require visual abilities. Since visual CAPTCHAs are still the most common type of CAPTCHAs, replacing them in practice will take time and require the creation of new standards. In the meanwhile, assistive transfer systems like WEBALLY can help improve web accessibility for PVIs more immediately.

While we focused on task-based CAPTCHAs in this work, the assistive transfer system approach can also be further explored to support other online tasks, for instance, helping PVI users screen images before they share them on social media to limit potential privacy leakage.

8 Conclusion

To help PVIs overcome task-based visual CAPTCHAs that frustrate and encumber their daily web use, we designed and implemented a proof-of-concept assistive transfer system — WEBALLY — that allows PVIs to source just-in-time, direct help from friends or trained crowd workers. Through an exploratory lab study with recruited PVIs and helpers, we found that both PVIs and helpers had a generally positive impression towards WEBALLY, finding it to be a useful alternative to other accessibility solutions (e.g., audio CAPTCHAs). Participants also found that WEBALLY offered sufficient mitigation to protect PVIs’ privacy and security in enabling limited remote control for task transfer. We also discovered several factors that may affect PVI participants’ perception towards using WEBALLY, such as the type of website in which the CAPTCHA is embedded, helper availability, and the potential impact such a system might have on a PVI and their helpers’ social relationships. Helpers, too, had varied preferences in terms of how frequently they would their help solicited and how they would want to be compensated for their effort.

In conclusion, assistive transfer systems like WEBALLY could serve as a preferred “last resort” alternative for PVIs when they cannot solve reCAPTCHA tasks independently. However, future work is needed to ensure timely recruitment of help (e.g., through proactive and speculative recruitment) and to establish compensation structures with which both PVIs and helpers feel comfortable. More broadly, we foresee assistive transfer systems as a promising new class of assistive technologies that can empower PVIs to overcome web accessibility hurdles related to security and privacy.

9 Acknowledgement

We thank our participants for overcoming difficulties to schedule meetings with several attendees and their insightful and articulated feedback, especially the help from Aditi Shah. We also thank the anonymous reviewers for their thoughtful comments and suggestions. This work was supported in part by the National Science Foundation (NSF Grant CNS-1652497).

References

- [1] Dina A Abdrabo, Tarek Gaber, and M. Wahied. Assistive technology solution for blind users based on friend-sourcing. In *The 1st International Conference on Advanced Intelligent System and Informatics (AIS2015), November 28-30, 2015, Beni Suef, Egypt*, pages 413–422. Springer, 2016.
- [2] Tousif Ahmed, Roberto Hoyle, Kay Connelly, David Crandall, and Apu Kapadia. Privacy concerns and behaviors of people with visual impairments. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 3523–3532, 2015.
- [3] Tousif Ahmed, Apu Kapadia, Venkatesh Potluri, and Manohar Swaminathan. Up to a limit? privacy concerns of bystanders and their willingness to share additional information with visually impaired users of assistive technologies. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(3):1–27, 2018.
- [4] Tousif Ahmed, Patrick Shaffer, Kay Connelly, David Crandall, and Apu Kapadia. Addressing physical safety, security, and privacy for people with visual impairments. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 341–354, 2016.
- [5] Taslima Akter, Tousif Ahmed, Apu Kapadia, and Swami Manohar Swaminathan. Privacy considerations of the visually impaired with camera based assistive technologies: Misrepresentation, impropriety, and fairness. In *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, pages 1–14, 2020.
- [6] Taslima Akter, Bryan Dosono, Tousif Ahmed, Apu Kapadia, and Bryan Semaan. "I am uncomfortable sharing what i can't see": Privacy concerns of the visually impaired with camera based assistive applications. In *29th USENIX Security Symposium (USENIX Security 20)*, 2020.
- [7] Anti-CAPTCHA. <https://anti-captcha.com/>. Accessed: 2021-05-23.
- [8] Paul Baecher, Niklas Büscher, Marc Fischlin, and Benjamin Milde. Breaking reCAPTCHA: a holistic approach via shape recognition. In *IFIP International Information Security Conference*, pages 56–67. Springer, 2011.
- [9] Daniel Robert Bateman, Erin Brady, David Wilkerson, Eun-Hye Yi, Yamini Karanam, and Christopher M Callahan. Comparing crowdsourcing and friendsourcing: a social media-based feasibility study to support alzheimer disease caregivers. *JMIR research protocols*, 6(4):e56, 2017.
- [10] BeMyEyes. <https://www.bemyeyes.com/>. Accessed: 2021-05-23.
- [11] Cynthia L Bennett, Erin Brady, and Stacy M Branham. Interdependence as a frame for assistive technology research and design. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility*, pages 161–173, 2018.
- [12] Michael S Bernstein, Joel Brandt, Robert C Miller, and David R Karger. Crowds in two seconds: Enabling real-time crowd-powered interfaces. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pages 33–42, 2011.
- [13] Michael S Bernstein, Desney Tan, Greg Smith, Mary Czerwinski, and Eric Horvitz. Personalization via friend-sourcing. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 17(2):1–28, 2008.
- [14] Rudy Berton, Ombretta Gaggi, Agnieszka Kolasinska, Claudio Enrico Palazzi, and Giacomo Quadrio. Are captchas preventing robotic intrusion or accessibility for impaired users? In *2020 IEEE 17th Annual Consumer Communications & Networking Conference (CCNC)*, pages 1–6. IEEE, 2020.
- [15] Jeffrey P Bigham and Anna C Cavender. Evaluating existing audio captchas and an interface optimized for non-visual use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1829–1838, 2009.
- [16] Jeffrey P. Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C. Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samuel White, and others. VizWiz: nearly real-time answers to visual questions. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*, pages 333–342. ACM, 2010.
- [17] Kevin Bock, Daven Patel, George Hughey, and Dave Levin. uncaptcha: a low-resource defeat of recaptcha's audio challenge. In *11th USENIX Workshop on Offensive Technologies (WOOT 17)*, 2017.

- [18] Richard E Boyatzis. *Transforming qualitative information: Thematic analysis and code development*. sage, 1998.
- [19] Erin Brady, Meredith Ringel Morris, and Jeffrey P Bigham. Friendsourcing for the greater good: perceptions of social microvolunteering. In *Second AAAI Conference on Human Computation and Crowdsourcing*, 2014.
- [20] Anjali Avinash Chandavale, Ashok M Sapkal, and Rajesh M Jalnekar. Algorithm to break visual captcha. In *2009 Second International Conference on Emerging Trends in Engineering & Technology*, pages 258–262. IEEE, 2009.
- [21] Geetest CAPTCHA demo. <https://www.geetest.com/en/demo>. Accessed: 2021-05-23.
- [22] Bryan Dosono, Jordan Hayes, and Yang Wang. “I’m Stuck!”: A Contextual Inquiry of People with Visual Impairments in Authentication. In *Symposium on Usable Privacy and Security (SOUPS)*, 2015.
- [23] Paul Dourish, Rebecca E Grinter, Jessica Delgado De La Flor, and Melissa Joseph. Security in the wild: user strategies for managing security as an everyday, practical problem. *Personal and Ubiquitous Computing*, 8(6):391–401, 2004.
- [24] Valerie Fanelle, Sepideh Karimi, Aditi Shah, Bharath Subramanian, and Sauvik Das. Blind and human: Exploring more usable audio captcha designs. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*, pages 111–125, 2020.
- [25] Buster: CAPTCHA Solver for Humans. <https://chrome.google.com/webstore/detail/buster-captcha-solver-for-mpbjkejclgfgadiemefgebjfooflflhl>. Accessed: 2021-05-23.
- [26] Jennifer LoCascio Gauvreau. Accessibility gotchas with captchas. *Journal of Digital & Social Media Marketing*, 5(4):379–390, 2017.
- [27] Danna Gurari, Qing Li, Chi Lin, Yanan Zhao, Anhong Guo, Abigale Stangl, and Jeffrey P Bigham. Vizwiz-priv: a dataset for recognizing the presence and purpose of private visual information in images taken by blind people. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 939–948, 2019.
- [28] Jordan Hayes, Smirity Kaushik, Charlotte Emily Price, and Yang Wang. Cooperative Privacy and Security: Learning from People with Visual Impairments and Their Allies. In *Proceedings of Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, 2019.
- [29] Kieran Healy. *Last best gifts: Altruism and the market for human blood and organs*. University of Chicago Press, 2010.
- [30] José María Gómez Hidalgo and Gonzalo Alvarez. Captchas: An artificial intelligence application to web security. In *Advances in Computers*, volume 83, pages 109–181. Elsevier, 2011.
- [31] J. Holman, J. Lazar, and J. Feng. Investigating the Security-related Challenges of Blind Users on the Web. In Patrick Langdon BSc, CEng John Clarkson MA MIEE, and Peter Robinson MA Ceng, editors, *Designing Inclusive Futures*, pages 129–138. Springer London, 2008.
- [32] Mohit Jain, Rohun Tripathi, Ishita Bhansali, and Pratyush Kumar. Automatic generation and evaluation of usable and secure audio recaptcha. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*, pages 355–366, 2019.
- [33] Kiranjot Kaur and Sunny Behal. Captcha and its techniques: a review. *International Journal of Computer Science and Information Technologies*, 5(5):6341–6344, 2014.
- [34] Sushama Kulkarni and Hanumant Fadewar. Audio captcha techniques: A review. In *Proceedings of the Second International Conference on Computational Intelligence and Informatics*, pages 359–368. Springer, 2018.
- [35] Jonathan Lazar, Jinjuan Feng, Tim Brooks, Genna Melamed, Brian Wentz, Jon Holman, Abiodun Olalere, and Nnanna Ekedebe. The soundsright captcha: an improved approach to audio human interaction proofs for blind users. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 2267–2276, 2012.
- [36] Il’dar Nailevich Manashev. Breaking google recaptcha v2. *Prikladnaya Diskretnaya Matematika. Supplement*, (11):99–101, 2018.
- [37] Lourdes Moreno, María González, and Paloma Martínez. Captcha and accessibility. *Is this the best we can do*, 2014.
- [38] Greg Mori and Jitendra Malik. Recognizing objects in adversarial clutter: Breaking a visual captcha. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 1, pages I–I. IEEE, 2003.
- [39] Meredith Ringel Morris, Kori Inkpen, and Gina Venolia. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, pages 662–673, 2014.

- [40] Meredith Ringel Morris, Jaime Teevan, and Katrina Panovich. What do people ask their social networks, and why? a survey study of status message q&a behavior. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1739–1748, 2010.
- [41] Donatella Pascolini and Silvio Paolo Mariotti. Global estimates of visual impairment: 2010. *The British Journal of Ophthalmology*, 96(5):614–618, May 2012.
- [42] Eyecoming Platform. <https://www.eyecoming.com/english/>. Accessed: 2021-05-23.
- [43] Frances L Rapport and CJ Maggs. Titmuss and the gift relationship: altruism revisited. *Journal of advanced nursing*, 40(5):495–503, 2002.
- [44] JAWS Screen Reader. <https://www.freedomscientific.com/products/software/jaws/>. Accessed: 2021-05-23.
- [45] NVDA Screen Reader. <https://www.nvaccess.org/>. Accessed: 2021-05-23.
- [46] Google reCAPTCHA. <https://www.google.com/recaptcha/about/>. Accessed: 2021-05-23.
- [47] Jeffrey M Rzeszotarski and Meredith Ringel Morris. Estimating the social costs of friendsourcing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2735–2744, 2014.
- [48] Stuart E Schechter, Rachna Dhamija, Andy Ozment, and Ian Fischer. The emperor’s new security indicators. In *2007 IEEE Symposium on Security and Privacy (SP’07)*, pages 51–65. IEEE, 2007.
- [49] Sajad Shirali-Shahreza and M Hassan Shirali-Shahreza. Accessibility of captcha methods. In *Proceedings of the 4th ACM workshop on Security and artificial intelligence*, pages 109–110, 2011.
- [50] Suphannee Sivakorn, Jason Polakis, and Angelos D Keromytis. I’m not a human: Breaking the google recaptcha. *Black Hat*, pages 1–12, 2016.
- [51] Online Survey. Webaim: screen reader user survey #7 results.
- [52] Jennifer Tam, Jiri Simsa, David Huggins-Daines, Luis Von Ahn, and Manuel Blum. Improving audio captchas. In *Symposium On Usable Privacy and Security (SOUPS)*, 2008.
- [53] Kathleen D Vohs, Nicole L Mead, and Miranda R Goode. The psychological consequences of money. *science*, 314(5802):1154–1156, 2006.
- [54] Michael Walzer. *Spheres of justice: A defense of pluralism and equality*. Basic books, 2008.
- [55] Yang Wang. Inclusive Security and Privacy. *IEEE Security & Privacy*, 16(4):82–87, August 2018.
- [56] Webvisum. <https://www.webvisum.com/en/main/download>. Accessed: 2021-05-23.
- [57] Jacob O Wobbrock, Shaun K Kane, Krzysztof Z Gajos, Susumu Harada, and Jon Froehlich. Ability-based design: Concept, principles and examples. *ACM Transactions on Accessible Computing (TACCESS)*, 3(3):1–27, 2011.
- [58] Shaomei Wu and Lada A Adamic. Visually impaired users on an online social network. In *Proceedings of the sigchi conference on human factors in computing systems*, pages 3133–3142, 2014.
- [59] Huichuan Xia, Yang Wang, Yun Huang, and Anuj Shah. "Our Privacy Needs to Be Protected at All Costs": Crowd Workers’ Privacy Experiences on Amazon Mechanical Turk. *Proc. ACM Hum.-Comput. Interact.*, 1(CSCW):113:1–113:22, December 2017.
- [60] Jeff Yan and Ahmad Salah El Ahmad. Breaking visual captchas with naive pattern recognition algorithms. In *Twenty-Third Annual Computer Security Applications Conference (ACSAC 2007)*, pages 279–291. IEEE, 2007.
- [61] Yuan Zhou, Zesun Yang, Chenxu Wang, and Matthew Boutell. Breaking google recaptcha v2. *Journal of Computing Sciences in Colleges*, 34(1):126–136, 2018.
- [62] Haiyi Zhu, Sauvik Das, Yiqun Cao, Shuang Yu, Aniket Kittur, and Robert Kraut. A market in your social network: The effects of extrinsic rewards on friendsourcing and relationships. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 598–609, 2016.
- [63] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. Research through design as a method for interaction design research in hci. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 493–502, 2007.

A Appendix

A.1 Pre-Study Interview (Requester Only)

- **I know what a CAPTCHA is and have encountered it in my real-life prior to this study**
 - o Yes
 - o No
- **Roughly, how many times did you need to solve a CAPTCHA in the past 7 days**
- **Roughly, how many times did you need to solve a CAPTCHA in the past 30 days**
- **Where did you usually encounter the CAPTCHA tasks**
- **I am confident in solving a Google reCAPTCHA vision task on my own (the task that ask users to click the right tile images)**
 - o Strongly Agree
 - o Agree
 - o Neither Agree nor Disagree
 - o Disagree
 - o Strongly Disagree
- **I am confident in solving a Google reCAPTCHA audio task on my own (the task that ask users to type in words they hear)**
 - o Strongly Agree
 - o Agree
 - o Neither Agree nor Disagree
 - o Disagree
 - o Strongly Disagree
- **Can you briefly describe why are you confident or not in solving these tasks?)**
- **Please explain your obstacles/challenges when solving a CAPTCHA**
- **When I need help for solving a CAPTCHA, I am confident that someone will always be able to help me in a timely fashion**
 - o Strongly Agree
 - o Agree
 - o Neither Agree nor Disagree
 - o Disagree

o Strongly Disagree

A.2 Post-Study Interview

A.2.1 Requester Interview

- **In general, what do you think about the tool?**
- **What was good or bad? What could be improved?**
- **Do you have any concerns using the tool? Can you explain them briefly?**
- **We saw that you chose [some methods] in solving the CAPTCHA tasks, Could you tell us why you chose the method(s)?**
- **If you have the choices, which option(s) would you choose to solve CAPTCHAs in the future? (Please explain why)**
 - o Free friend-sourced tool
 - o Free crowd-sourced tool
 - o Paid friend-sourced tool
 - o Paid crowd-sourced tool
 - o I'd not choose any of the above options
- **I am confident in solving a CAPTCHA with free, friend-sourcing using the tool (Please explain**
 - o Strongly Agree
 - o Agree
 - o Neither Agree nor Disagree
 - o Disagree
 - o Strongly Disagree
- **I am willing to install and use the tool**
 - o Strongly Agree
 - o Agree
 - o Neither Agree nor Disagree
 - o Disagree
 - o Strongly Disagree
- **I am willing to recommend the tool to others**
 - o Strongly Agree
 - o Agree
 - o Neither Agree nor Disagree
 - o Disagree
 - o Strongly Disagree
- **The tool will bring me and my helping friend closer**
 - o Strongly Agree
 - o Agree

- Neither Agree nor Disagree
- Disagree
- Strongly Disagree

- **I have privacy or security concerns when using the friend-sourced tool**

- Strongly Agree
- Agree
- Neither Agree nor Disagree
- Disagree
- Strongly Disagree

- **I have privacy or security concerns when using the crowd-sourced tool**

- Strongly Agree
- Agree
- Neither Agree nor Disagree
- Disagree
- Strongly Disagree

- **Do you have any other suggestions for improving this tool?**

A.2.2 Helper Interview

- **In general, what do you think about the tool?**

- **What was good or bad? What could be improved?**

- **Do you have any concerns using the tool? Can you explain them briefly?**

- **If you have the choices, which option(s) would you choose to help solve CAPTCHAs in the future? (Please explain why)**

- Helping your friend/family member without getting paid
- Helping your friend/family member and getting paid for a small amount
- Helping strangers without getting paid
- Helping your strangers and getting paid for a small amount
- I'd not choose any of the above options

- **I am willing to be prompted and help the requester (Please explain)** Strongly Agree

- Agree
- Neither Agree nor Disagree
- Disagree
- Strongly Disagree

- **I think this tool will be too time-consuming (Please explain)** Strongly Agree

- Agree
- Neither Agree nor Disagree
- Disagree
- Strongly Disagree

- **If requesters request it, I can always respond timely**

- Strongly Agree
- Agree
- Neither Agree nor Disagree
- Disagree
- Strongly Disagree

- **The tool will bring me and my helping friend closer**

- Strongly Agree
- Agree
- Neither Agree nor Disagree
- Disagree
- Strongly Disagree

- **I have privacy or security concerns when helping friends** Strongly Agree

- Agree
- Neither Agree nor Disagree
- Disagree
- Strongly Disagree

- **I have privacy or security concerns when helping strangers** Strongly Agree

- Agree
- Neither Agree nor Disagree
- Disagree
- Strongly Disagree

- **Do you have any other suggestions for improving this tool?**