



ImageAlly: A Human-AI Hybrid Approach to Support Blind People in Detecting and Redacting Private Image Content

Zhuohao (Jerry) Zhang, University of Washington, Seattle; Smirity Kaushik and JooYoung Seo, University of Illinois at Urbana-Champaign; Haolin Yuan, Johns Hopkins University; Sauvik Das, Carnegie Mellon University; Leah Findlater, University of Washington, Seattle; Danna Gurari, University of Colorado Boulder; Abigale Stangl, University of Washington, Seattle; Yang Wang, University of Illinois at Urbana-Champaign

<https://www.usenix.org/conference/soups2023/presentation/zhang>

This paper is included in the Proceedings of the Nineteenth Symposium on Usable Privacy and Security.

August 7-8, 2023 • Anaheim, CA, USA

978-1-939133-36-6

Open access to the Proceedings of the Nineteenth Symposium on Usable Privacy and Security is sponsored by USENIX.

ImageAlly: A Human-AI Hybrid Approach to Support Blind People in Detecting and Redacting Private Image Content

Zhuohao (Jerry) Zhang¹, Smirity Kaushik², JooYoung Seo², Haolin Yuan³
Sauvik Das⁴, Leah Findlater¹, Danna Gurari⁵, Abigale Stangl¹, Yang Wang²

¹University of Washington, Seattle ²University of Illinois at Urbana-Champaign

³John Hopkins University ⁴Carnegie Mellon University ⁵University of Colorado Boulder

{zhuohao, Leahkf, astangl}@uw.edu}@uw.edu, {smirity2, jseo1005, yvw}@illinois.edu

{hyuan4}@jhu.edu, {sauvik}@cmu.edu, {danna.gurari}@colorado.edu

Abstract

Many people who are blind take and post photos to share about their lives and connect with others. Yet, current technology does not provide blind people with accessible ways to handle when private information is unintentionally captured in their images. To explore the technology design in supporting them with this task, we developed a design probe for blind people — ImageAlly — that employs a human-AI hybrid approach to detect and redact private image content. ImageAlly notifies users when potential private information is detected in their images, using computer vision, and enables them to transfer those images to trusted sighted allies to edit the private content. In an exploratory study with pairs of blind participants and their sighted allies, we found that blind people felt empowered by ImageAlly to prevent privacy leakage in sharing images on social media. They also found other benefits from using ImageAlly, such as potentially improving their relationship with allies and giving allies the awareness of the accessibility challenges they face.

1 Introduction

A challenge for blind people¹ is how to remove private information they unintentionally capture in images they take before sharing the content with others (e.g., personal information on stray screens or pieces of paper, human faces that were not supposed to appear). For example, prior work reported that over 10% of over 40,000 images taken by blind people contained private information [19]. Yet, sharing images is a key

¹We use the identity-first language when describing people with visual impairments, guided by the National Federation of the Blind.

way for people to connect with each other, including on social networking services (SNSs) [36]. This challenge on how to preserve private visual information is relevant for the more than 49 million blind people around the world [1, 11, 29, 50, 51]

Given the increasing ubiquity and accessibility of mobile devices with built-in cameras, there is a growing potential benefit of developing technology that supports blind users in redacting private information in images. Yet, this capability is not yet available. For instance, a potential workaround is to leverage existing image editing tools to redact private information in images, yet such tools are inaccessible to blind people. That is because such tools require precise hand-eye coordination (e.g., moving the mouse or finger to brush over specific areas). *Our goal is to bridge this gap by empowering blind people with an accessible tool that facilitates the detection and redaction of private content in images they intend to share with others.*

We introduce a new human-AI hybrid approach to enable blind people to avoid unintended privacy-violating disclosures in images they intend to share publicly. We first employ computer vision to provide first-pass prescriptive insights (object recognition and image captioning with associated confidence scores) about image content. Blind users can then decide for themselves whether they consider these identified objects as unnecessarily private or sensitive and whether they want to ask their trusted sighted allies (family members or friends) for targeted editing assistance. If they choose to do so, they can specify how they want the image to be edited (e.g., blurring specific human faces or cropping out certain parts of the image). Our approach was inspired, in part, by what Nissenbaum et al. call *handoff* [40] and what Zhang et al. call an *assistive transfer system* [54]: a system that allows blind people to solicit just-in-time, targeted assistance from a trusted sighted ally to solve an outstanding accessibility challenge. This human-AI hybrid approach also aligns with the emergent perspective of embracing interdependence in assistive technology design [8]. We designed a proof-of-concept system to operationalize and assess this hybrid approach: *ImageAlly*. We aim to use this system as a probe to understand how such tools can be used by blind users and their sighted allies [26].

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) 2023, August 6–8, 2023, Anaheim, CA, USA

We conducted a user study to deploy this probe with 20 participants (10 pairs of one blind individual and one sighted ally), following a pilot study (see appendix A.2) with seven participants (four blind people and three sighted allies). The goal of our study was to answer three research questions:

1. How well does our human-AI hybrid approach address blind people’s need to identify and redact private information in images they consider sharing online?
2. Given that both the AI-generated insights and the edits generated by human allies provide different levels of information and carry some uncertainty (e.g., AI results can be inaccurate and allies’ edits can be subjective), how might blind people use human versus AI assistance? As part of this, we are interested in how they might deal with the uncertainty when deciding whether to solicit targeted editing assistance from trusted allies and share edited images online.
3. How might use of ImageAlly affect the perceived relationship between blind people and their trusted allies, given that prior research on friendsourcing in general [56] and assistive transfer systems in particular [54] suggests that friendsourcing approaches can impact social relationships between friends?

We found that the ImageAlly approach showed promise in supporting blind people in sharing images in a way that aligns with their personal privacy preferences by facilitating the detection and redaction of private content in those images. We also found that our blind participants varied in what they wanted out of ImageAlly. For example, our participants wanted different things out of the AI-powered image screener: some preferred minimal descriptions of image content so they could *efficiently* check for privacy leaks in images they took themselves, while others wanted to use the screener to confirm that their allies appropriately redacted private information in their images. For images processed by allies, we also observed some inconsistencies between sighted allies’ editing and blind users’ preferences in six out of a total of 20 cases, which were perceived differently by different blind participants and could potentially be avoided by using ImageAlly for a second-time AI screening. Lastly, some participants also believed that their interactions with the blind individuals or sighted allies through ImageAlly have potentially positive impact on their relationships. For example, some sighted ally participants felt that ImageAlly has a positive value in improving their awareness of the challenges that their blind family members or friends faced. They also found ImageAlly useful and felt it could prevent their blind family members or friends from accidentally sharing private information with others.

To summarize, our work makes three main contributions: (1) we introduced and explored the design space of assistive transfer systems for processing images with private information, (2) we designed and implemented a proof-of-concept

assistive transfer system, ImageAlly, to serve as a design probe to explore our human-AI hybrid approach in facilitating the detection and redaction of private photo information for blind people, and (3) we conducted a design probe study with both blind people and their sighted allies to answer our research questions and synthesize design insights for assistive transfer systems and other tools designed to improve blind people’s exploration and editing of images.

2 Related Work

2.1 Image Sense-making for Blind People

One approach that blind people currently take when they want to make sense of images is to rely on apps and services that use state-of-the-art computer vision techniques to detect objects and caption images. Such tools include Microsoft’s Seeing AI [2], Google’s Lookout App [16] and other automated image description services [52]. While these services describe image content, their outputs do not offer prescriptive guidance to assist blind people in identifying and obfuscating private information that may be unintentionally captured in those images.

Besides commercially available tools, the development of deep learning models [17, 31, 42] has spurred the increasing use of automated image description models that can assist in image sense-making [5, 25, 53]. Such technologies simplify a wide range of everyday tasks including identifying objects and recognizing familiar faces or facial expressions [4]. Zhao et al. studied how state-of-the-art computer-generated descriptions in Facebook’s photo-sharing feature can help blind people improve the photo-sharing experience [55]. Blind people were also found to place a lot of trust in automatically generated captions for visual content on social media (e.g., Twitter) although the caption may diverge from the visual content [37]. Simons et al. studied crowd workers’ motivations and challenges for generating image descriptions to develop automated solutions [44]. Finally, Gurari et al. explored the limitations of modern algorithms in captioning images taken by blind people [20]. While ImageAlly is guided by these prior studies, our design probe is novel in understanding blind users’ practices of handling private visual content in a human-AI hybrid fashion.

Another approach blind people currently employ to make sense of images is to rely on human intelligence through crowdsourcing or friendsourcing. Blind people sometimes solicit sighted assistance from remote humans to support visual interpretation and visual question answering tasks. This includes relying on remote professional assistance services, such as Aira [3], asking physically proximate allies for direct assistance [51], and soliciting assistance using commercial and research-based crowdsourcing services, including Be My Eyes [41] and VizWiz [10]. Generally, these services provide blind people with remote assistance from sighted allies who, for example, answer questions about their surroundings or provide vocal-guidance on using inaccessible interfaces. However, a limitation of these human-based services is that they do not directly support screening images (or videos) for private content.

2.2 Photo Practices of Blind People

Blind individuals (including teens [9]) take and share photographs for the same reasons that sighted people do [6, 22, 29]. However, when they share photos, identifying possible private and sensitive information inside can be a challenge [45], not to mention processing that information. Researchers have developed many alternatives to assist blind photography. For example, Google Lookout App [16] and iOS AI-based VoiceOver recognition [49] provide text and audio feedback of what is in the camera field of view.

Vázquez et al. proposed to help blind users aim a camera so that they can know for sure what content is inside the frame [47]. Iwamura et al. [28] tried to solve the same problem by introducing a system that uses an omnidirectional camera. Complementing the plethora of prior work around blind photography, we introduce and evaluate the first prototype directly designed to empower blind photographers to avoid inadvertently sharing private/sensitive content captured in their images. Our work builds off of prior work that investigated obfuscation techniques to mitigate privacy leakage in images, such as via blurring, pixelating, inpainting, and avatars [27, 35]. Our ImageAlly system provides sighted allies with obfuscation options and, to our knowledge, our work is the first to explore blind people’s experiences of applying obfuscation for privacy-preservation.

3 Design Considerations

We began designing ImageAlly by identifying the challenges blind people experience with sharing images, especially when those images might contain private information. More generally, we identified three design goals for a hybrid human-AI system for blind users to identify and occlude private information in images based on recommendations from prior literature [45, 55]. A co-author of this paper who is blind also informed the design goals we strove towards.

3.1 Identifying Potentially Private and Sensitive Information in Images

Our first goal is to fine-tune state-of-the-art computer vision models to identify potentially private information in an image to facilitate targeted editing or further description by a sighted ally. This goal emerged based on our understanding that despite impressive advances in facial recognition, object detection, optical character recognition (OCR), and image captioning, the identification of what content may be private is highly contextual and personal [4, 45]. Thus, full delegation of this responsibility to AI may be untenable. Such a system would require going beyond simply identifying and captioning the image contents, to recognizing the content in relation to blind people’s specific visual privacy concerns in a particular sharing context. In turn, our objective was to fine-tune existing AI models to better identify image content that blind people

might consider to be private, e.g. faces and text [45, 55], and then transfer that image content to sighted human allies who can further interpret and redact the private content. We consider such information screening as a first-pass prescriptive insight.

3.2 Redacting Private Information in Images

Our second goal is to source human assistance to redact private information and provide description of their operations for blind people to digest what has been changed in images. For example, blind people may want to crop a certain application window out of a screenshot of their laptop screen, blur out personally identifiable information in a document scan, or blur out children’s faces in personal photos. These tasks require an accurate understanding of private visual information and precise hand-eye coordination to act on it, such as moving the mouse/finger to the edge of or over the area that needs to be blurred. Therefore, rather than trying to build an automatic image editing tool using AI models, our objective was to source direct human assistance to help edit the photos. One opportunity to address this goal is to enable a remote ally to directly edit the photos. We consider such a photo transferring and editing as a second-pass human-powered editing.

3.3 Communicating and Verifying Screening Preferences

Our third goal is more of an additional consideration to complement the human-AI hybrid approach. While such a hybrid system of (1) first-pass AI-generated prescriptive insights and (2) second-pass human editing can address our first two design goals, privacy needs vary across individuals and sharing scenarios. For example, the designated ally might recognize some personally identifiable information as private, and not recognize other information as private, and return a photo that does not meet the blind user’s expectation (e.g., blurring faces that were not meant to be blurred, or forgetting to crop parts of the image that were meant to be redacted). Given that blind people may not be able to confirm if the edited photo was edited in line with their expectations, a third goal was to provide an accessible way for blind people and their allies to communicate expectations, preferences, and actions. Blind people should have a way to directly state how they expect the photo to be edited.

4 ImageAlly System

Guided by our design goals, we implemented ImageAlly as a design probe [26] on iOS using React Native. In addition to the mobile app, ImageAlly also includes the interface for allies and the backend server. The interface for allies presents the photo sent for redacting, the blind requester’s instructions for how they would like the image edited, and an interactive image-editing tool. Next, we describe ImageAlly’s interfaces for the blind users and their allies.

4.1 Non-Visual Interface

We designed and developed ImageAlly’s non-visual interface (see figure 3) to provide blind users with image descriptions (descriptive screening results) through text (and sound when accessed through a screen reader). To do so, ImageAlly first employed existing libraries and APIs [43] to detect potentially privacy-intrusive information — i.e., faces, pre-selected object categories (e.g., documents, ID cards), and texts that appear in photos based on insights from prior work [19,45] and our blind coauthor’s personal experience. This variety of information was collected to assist blind users in their decision-making process while having them ultimately determine if the identified content is private or sensitive, and if so, whether to edit it or leave it as is. Accordingly, in the case that the descriptive screening results indicate that there is potential private information in the image, the interface provides users with a choice to send the photo to their designated ally. As part of this process, the user can specify preferences—by choosing from a list of common options or by typing in their own preferred message for the ally—to indicate how the image is further evaluated for private information. Lastly, users will be asked to select a contact and click a button to send the photo-processing request. Figure 2 summarizes the interaction workflow of ImageAlly. We provide a detailed description of ImageAlly’s descriptive screening process in the appendix.

4.2 Visual Interface for The Allies

Once a blind user obtains AI-generated descriptions of potentially private content in images, they may next choose to solicit assistance from allies to redact this information. These sighted allies are solicited through an SMS or email message in which they are provided with a link. The link, in turn, directs the ally to a web interface that presents the photo to be edited, the blind users’ corresponding instructions for what information to redact, and a suite of controls to help with redacting private information in images. For example, if the blind user asks the ally to blur out all the text in the photo, the ally can use the built-in tools to blur out the image partly and return the image back to the blind person. Allies also have a text-input box through which they can inform the blind requester of what they did to the photo.

We provided as obfuscation techniques pixelating and blurring using finger-drawing (like an eraser). Note that we use the term “blur” in ImageAlly and throughout the paper as a general term for obfuscation, unless noted otherwise, since we used this term with our study participants to make it easier to understand than with a more technical term such as obfuscation. Of note, prior work has shown that obfuscation techniques such as blurring and pixelating can be ineffective [33–35, 48] or attacked (reversed) via deep learning [39], however, they are still favored by users and viewers [13, 23, 35, 48]. Considering the privacy-utility trade-off and the required effort for obfus-

cation, we chose *blurring* and *pixelating* in our current design as simple interactions for sighted allies to perform. With that said, ImageAlly could incorporate and work with other current and future improved obfuscation techniques.

5 Study Method

We used ImageAlly as a probe that serves the design goal of inspiring users and researchers to think about new technologies and the social science goal of understanding the needs and desires of users in a real-world setting [26]. Specifically, we sought to gain insights into blind users’ perceptions of, preferences towards, and usage of a hybrid human-AI assistive transfer system for identifying and redacting private information in photos they intend to share online. To that end, we conducted an exploratory study of ImageAlly with 20 participants (10 blind people and 10 allies) using an IRB-approved protocol. We asked blind people and one of their sighted allies (a friend or family member, recruited with the blind participant) to use ImageAlly to screen and edit photos from different sources.

After the study, we conducted a comparative analysis on the pictures initially selected by the blind participants and the redacted pictures edited by the allies. This comparative analysis highlighted differences in how participants used ImageAlly, as well as afforded us insight into whether ally-edits aligned with blind participants’ preferences. Furthermore, for those ally-edited pictures that did not fully match blind participants’ preferences, we followed up with the blind participants and asked them how they felt about and wanted to act on that inconsistency. Together we evaluated how ImageAlly worked in detecting and redacting private image contents and covered cases where ImageAlly did not work perfectly and how blind users would like to handle it.

5.1 Participants

We recruited participants in pairs: one blind user and one sighted ally who the blind user considered a trusted friend or family member. In total, we recruited 10 pairs for 20 total participants: 10 blind participants (referred to as requesters and numbered from R1 to R10), and 10 sighted allies accordingly (referred to as A1 to A10). Also two requesters reported to have hearing impairments. The relationship of the participant pairs varied from friends to family members including mother and daughter, brother and sister, husband and wife. Note that we only recruited requesters that use iPhones.

5.2 Apparatus

We used the ImageAlly design probe to conduct the exploratory lab study. We provided a downloadable link via TestFlight [46] before the study to let the blind users install ImageAlly on their iPhone. For sighted allies, we also designed a simple web interface that contains basic image editing tools including

functions, such as blurring, cropping, and drawing overlay markups (Figure 3, right). We developed the system so sighted allies would not themselves need to install ImageAlly, but would instead receive SMS text or email messages with an embedded link assigned to this photo-editing session.

To simulate different scenarios, we used two image sources in the study. First, we asked each blind user to prepare an image that contains information they consider private. To protect their privacy, we asked them to use outdated information: (1) their room surroundings, (2) selfie or family photos, (3) a screenshot of phone chat history, (4) a received letter, (5) an expired credit card, (6) an expired ID card, (7) a medicine bottle with descriptions, (8) visited webpages. These are the main privacy categories identified in the VizWiz-Priv dataset [19]. We ensured using their photos only for this project.

Second, we asked the blind users to share/re-post an image prepared by our research team on social media. This is to simulate the situation where they share others' visual content. The image was a mobile phone screenshot of a work group's chat history with co-workers' names and avatar profiles.

By using two different sources, we were better able to evaluate ImageAlly by accounting for a broader variety of real world scenarios in which a blind person may consider soliciting assistance, e.g., capture photos, and/or share and repost photos from a second party.

5.3 Procedure

First, we conducted a single session remote study over Zoom. The remote aspect of the study enabled our research team to simulate the likely use-case for ImageAlly, where the requester and the ally are not co-located when the requester might need to use ImageAlly. Upon receiving participants' written consent, we video-recorded all sessions and took detailed notes.

All study sessions lasted about an hour, including a post-study interview to gain insights of requesters' and allies' feedback separately in Zoom breakout rooms. Prior to the study, the participants were told to prepare one photo from the categories mentioned above in 5.1.2, and install the ImageAlly App, which took around 5-10 minutes.

After the study preparation and introduction, the researchers divided the participants into two Zoom 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 each of the two breakout rooms to guide them through the study, answer their questions, and conduct the exit interviews. 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 (shown in Appendix A.6) about their experience with photo sharing. The researcher in the ally's room (referred to as Researcher 2) also introduced the tasks and asked the allies questions about their previous experience

of receiving requests and providing visual assistance by describing the content of the photos, their concerns about seeing private information from others, and their preference about being contacted by requesters.

After asking both participants about their previous experiences with requesting and/or providing visual assistance, the researchers explained the possible scenarios in which ImageAlly could be used. We asked the requesters to imagine that they were sharing photos across two scenarios that were meant to approximate distinct real-world situations in which blind people may want to share a photo but may harbor concerns about photo content: sharing original photos taken by themselves, and (re-)sharing photos taken by others. Doing so allows us to compare/contrast preferences across different contexts of use — for example, would participants have different privacy concerns when sharing others' vs. their own photos? Would participants want the system configured, and if so, how? To strengthen ecological validity from the ally's perspective, they were instructed that the requests from their friends may come at any time, and that they could do other tasks rather than passively waiting for ImageAlly requests. When their assistance was requested, they would be notified via SMS text message.

After explaining the scenarios, the lab study began. Researcher 1 asked the requester to navigate to the ImageAlly App, go over the instructions in the App, and follow all the prompts from step one to step four (figure 3 left). For the first session, Researcher 1 asked the requesters to use their own photos and answered any question they may have during use of ImageAlly. Within the app, users have the option to send images to allies for editing depending on whether they believe there was private information in images. However, in our study, because the blind participants were instructed to prepare photos with private information prior to the study, most of them (9 out of 10) chose to continue sending the request since there was private information in the images. Only one participant (R1) prepared a selfie which didn't have private information they wanted to blur. However, R1 also chose to continue exploring the full features of ImageAlly. Before they sent the request, the blind participants were prompted to select a contact from their contact book integrated inside ImageAlly. Then they clicked a button to send the request.

After the allies received the message, they were asked to open a link with the web interface of the image editing tool inside. They were also prompted to follow the requesters' stated preferences and crop or blur out certain parts of the photos and provide text description of what they did to the photo. After they were finished, they clicked a button to send back the photo and the description, which can be saved and read by the requesters.

The session was repeated for the second scenario where requesters were asked to forward a photo created by our research team. The photo was sent to the requesters either using an email attachment or a Dropbox download link. Their allies also followed the same process of receiving, reviewing, and editing images. After completing all the tasks,

the researchers conducted an exit interview asking about participants' detailed experiences with the ImageAlly system. The questions include general feedback, suggestions for improvement, Likert-scale questions on system usability, and how such requests might impact the social relationship dynamics between requesters and allies (see Appendix A.7).

After the study sessions, we observed that for some images, there were some inconsistencies between the blind participants' preferences, sighted allies' edited images, and/or their descriptions of how they edited the images. We then analyzed those images and further followed up with the blind participants whose preferences were not fully addressed in the edited images. We asked those participants how they felt about and how they would handle the inconsistency. We report on this post-study analysis in Section 6.4.

5.4 Data Analysis

Upon receiving participant consent, we recorded the study Zoom meeting and logged users' behaviors in the ImageAlly prototype (e.g., blind users' requests and preferences, and how allies edited images) as suggested by Hutchinson et al. [26]. We transcribed the study videos and two members of the research team analyzed the study sessions using thematic analysis [12]. We first individually read and familiarized ourselves with the transcripts. Next, we performed an open coding of sessions independently. We then discussed regularly and eventually converged on the codes and the groupings of codes (i.e., themes) emerged. Since this work is exploratory and our analysis involved the generation of new codes, following guidelines from prior work [38], we did not calculate inter-coder reliability. Example themes include blind participants' prior experience, general feelings about ImageAlly, and how they reacted to the AI and ally generated results about photos. For sighted allies, our analysis covered their impressions about ImageAlly, their willingness and general availability to help blind requesters with images. We also recorded task completion time and their responses to System Usability Scores [15].

We conducted a comparative analysis on the 20 pictures used in the study (10 original pictures selected directly by blind participants, and 10 selected by researchers to be "forwarded" by participants) and the processed version of those pictures edited by allies during the study. Two researchers examined each image manually by independently recording the differences between the picture sent by the blind participant and the picture edited by the ally. We mainly analyzed: (1) What was the image about? (2) What were the blind participant's preferences, how does the edited image look like, and what was the ally's description of how he or she edited the image? (3) What was the difference between the AI screening results of the original image and the edited image?

Then two researchers met online and discussed consensus of these recorded differences. These differences were fairly straightforward to annotate (e.g., whether a person's face

has been blurred or not). We mainly focused on gauging the alignment between blind participants' stated preferences for edits and how sighted allies actually edited the pictures.

6 Results

To contextualize our study findings, We first present results that answer our three research questions (from subsection 6.1 to 6.3 accordingly). Finally, we talk about a post-analysis of inconsistent image editing from allies (subsection 6.4).

6.1 Overall Impression and Use of ImageAlly

To answer our first research question about how blind participants and sighted allies felt about ImageAlly's features, we asked participants about their overall impression of ImageAlly. All participants reported that they liked the ImageAlly system but also identified specific pros and cons. To frame their reactions to the system, we next report on our observations of how participants used ImageAlly in the study.

6.1.1 ImageAlly Usage

All 10 blind participants successfully installed the App prior to the study or with a researcher's help during the study, selected the photos they wanted to screen, and received the AI screening results. The photos chose by participants for the study included selfies, family photos, screenshots with personally identifiable information, expired ID cards and credit cards, and document scans. We present details of these photos in table 2 (Appendix A.4).

Each task, from opening the app to receiving the edited photo and saving or sharing it, took between 5 to 15 minutes. We note that the allies were ready to help immediately, which might not always be the case in practice. Task duration depended on several factors such as how familiar both requesters and their allies were with the ImageAlly interface (for the second scenario/task) and whether they asked questions during the task. However, ImageAlly provides an asynchronous way to process photos and it is often not a time-sensitive or urgent task. Most blind participants (9 out of 10) said they were willing to wait for the request to be completed, since they understood that, for instance, the human-editing process could take time. Unlike the first photo task where blind participants used their own images, in the second photo task we asked participants to imagine that they were "*forwarding*" photos from others. All blind participants were able to follow the same steps as the first task.

We also discovered that blind participants used ImageAlly differently across the two scenarios. For their own photos, two blind participants (R2 and R8) took the photos days before the study and thus couldn't locate the photos instantly. Then they used ImageAlly's AI screening function as a confirmation tool to help find the right photo. R2 mentioned that ImageAlly provides a quick and accessible way to confirm whether this photo was the one they wanted to select because it's "*just a couple clicks away*." They already knew roughly what was in

the photo and could quickly recognize the simple keywords or key objects in the AI screening results. For example, R2 found the keyword “*server*” in the AI screening results and immediately recognized that was the right photo she intended to select. In contrast, when “forwarding” others’ images, most blind participants used ImageAlly as an exploration tool (as opposed to confirmation) in order to understand the contents of the images because they had no idea.

6.1.2 Blind Participants’ Impressions of ImageAlly

All blind participants liked ImageAlly’s overall functionality and workflow. They spoke positively about ImageAlly providing: (1) some level of independence and interdependence (e.g., R2 said “*I love the idea, it basically gives me the freedom to do stuff myself, and it’s a really great way for my family to assist me*”); (2) more information about photos (e.g., R4 said “*I think it’s really cool. I have a lot of experiences with image description but all of them are limited. It’s giving me much information*”); and (3) accessible and user-friendly interface (e.g., R3 commented “*I enjoyed it because it’s simple to use. It makes sense once I get it and it’s pretty user-friendly*”).

Many participants (five blind participants and six ally participants) also pointed out the limitations of ImageAlly. First, while the system allows allies to provide a description of what the photo is about and what they did with the photo, it was still sometimes hard for blind requesters to know what was changed and to trust that the edited photo was free of private information. Two blind participants (out of 10) expressed concern about the edited photos. For instance, R1 mentioned that “*Maybe they missed something or I missed something. Before I share, I want to be confident of what to share.*” R1 was worried that perhaps her preference recorded wasn’t clear enough for the ally or the ally misinterpreted the message and edited it unexpectedly. However, the other blind participants (8 out of 10) expressed their trust towards their friends and family members in whether they could successfully edit the photo as requested. For instance, R3 said that “*If I choose this friend to send the image, it means I trust them and along with the photo they edited*”. R2 had a similar sentiment: “*I don’t need another way (to confirm) because I trust my friend and family*” Note that ImageAlly does not introduce a new trust challenge — even with face-to-face assistance, blind people still face the same challenges with trusting that their ally accurately edited their photo in accordance with their preference. In fact, ImageAlly provides a partial solution to this trust challenge: requesters can run an edited image through the ImageAlly to get some descriptive insight into how an ally’s edits changed what was perceivable to the AI, and can just as easily solicit a second opinion from another trusted ally.

Participants also suggested other areas for improvement. For instance, four blind participants said that they wanted to receive notifications when their allies received the request for photo editing to remove the private content, when the allies start working on the request, and when the allies finish

checking the photos. They also desired a way to check their allies’ availability before they send the request and the option of sending requests to multiple people at the same time when they are unsure if someone is not available.

6.1.3 Allies’ Impressions of ImageAlly

All ally participants found the tool useful and easy to use in general. Allies highlighted that ImageAlly could prevent their blind friends or family members from accidentally sharing sensitive information such as credit card information with others. They also felt that the tool readily provides them with ways to help their blind friends or family members. For instance, A8 shared that using ImageAlly would make her “*feel more confident when my husband has to send pieces of info to someone.*” She further highlighted that ImageAlly eliminates the need for her to be physically present to help her husband, “*He doesn’t necessarily need me right there, he can be in office and me at home and still help him out.*”

Some ally participants were even interested in using the blur feature of the tool for their own photos because they were not aware of any other tool that provides the similar blur feature. Our ally participants also offered design suggestions for the tool, such as the ability to zoom into the picture to precisely blur required information. Some allies (A2, A3, A8, A9) also found the instructions provided by the blind requester a bit confusing and hard to interpret. For instance, according to A8, “*Blur my identifiable information is confusing whether it should include only their information or everyone else’s too.*” Future designs of tools like ImageAlly could allow back-and-forth communication between requesters and helpers.

6.1.4 Perceived Usability of ImageAlly

We also used the System Usability Scale [15] to measure our participants’ perceived usability of ImageAlly. Most participants (both requesters and allies) agreed or strongly agreed that ImageAlly was easy to use, had well-integrated and consistent features, and that they would like to use it if ImageAlly is available (figure 1). The calculated SUS scores [15] were 86.25 for blind participants and 84.25 for sighted allies, indicating high usability of ImageAlly.

6.1.5 Privacy Concerns

ImageAlly was designed to process photos with private or sensitive information when blind people wish to share them, either on social media or with friends. However, people may have concerns even when sending them to friends or family members to check. Therefore, we explicitly asked about requesters’ concerns of sending photos to an ally. Our blind participants expressed that since ImageAlly allows them to choose the photo and the people they trust to ask for help, they were not concerned. For instance, R6 said “*Now I am confident of what I am going to share.*”. Similarly, ally




Figure 1: System Usability Scale scores from requesters and allies showing that both blind participants and their allies had a general positive attitude towards ImageAlly’s usability.

participants also reported no concerns seeing photos shared by blind participants. However, in some cases, the relationship might impact what photos a blind participant would share with the allies. For example, when A9 was asked, if she had any concerns seeing her mother’s private information in the photos, she responded, *“Not an issue so far. My dad is sighted he also checks photos with her.”*

6.2 AI and Human-Generated Results

Our second research question focused on how the blind participants perceive and use AI screening and human-processed results. For example, what AI-generated information would be useful for blind people to decide whether to share photos; and how they interpret and use this different information (e.g., type of objects identified, confidence scores of AI results).

6.2.1 Usage of AI-generated Results

Face Number Recognition. ImageAlly provides the number of faces detected in the photo, a feature that most of our blind participants found to be simple and effective. For instance, R8 said that *“When I take objects, I want to know if there are faces I am not aware of, this is quite important.”* R3 compared the simple face number recognition with SeeingAI’s image exploration functions and thought such detailed and thorough exploration of images in SeeingAI was not necessary when they are taking and checking photos. R5 also pointed out that when they are taking their own photos, they usually *“have a clue of what’s going on”* in the photo, and thus simple feedback such as the number of faces recognized is sufficient and more efficient than more detailed descriptors.

Object Detection and Text Extraction. In our study, eight out of 10 photos used by our blind participants had text-based private or sensitive information such as names, addresses, ID numbers (see table 2 in Appendix A.4 for details). All blind participants found that text extraction was particularly important when deciding whether to share photos. In practice, AI-generated outputs have inevitable uncertainty. We were

interested in how blind participants would make sense of and make use of the confidence scores of the privacy-relevant AI-generated outputs. For example, how would a confidence score of 50% versus 80% alter a blind user’s perception of and trust in the output? We noticed several occasions in the study where the confidence score of a certain object in the blind participant’s own photos was low; when that happened, we asked additional questions in the exit interviews about their interpretations of those outputs. We found that while high confidence scores on certain objects (e.g., text of addresses and ID numbers) would unsurprisingly motivate blind people to pay more attention to the photo and transfer the photos to their allies, low confidence scores of any object might also have a similar impact on our blind participants. For instance, R3 said that *“If the accuracy is low, it must be complex in the photos, so I would need it (ImageAlly) even more.”* In comparison, R4 interpreted a low score as something unacceptable, *“Low score doesn’t make sense to me. They are not as helpful so I would hesitate and even retake the photo before sending it to friends.”* Participants tend to believe that low confidence scores often represent complex situations in photos and thus they are more motivated to send requests to friends or family members to process the photos. However, it is hard to define a universal low-score threshold for everyone; standards may vary across individuals and contexts. While unpacking the effects of confidence score beyond the scope of our research, understanding uncertainty in AI-generated outputs for systems like ImageAlly appears to be a ripe area for future research.

6.2.2 Usage of Friendsourcing Results

After allies finished editing the photos and sent them back with descriptions, we asked requesters to read through the descriptions and to go over the sharing function either directly with contacts or on social media. Some participants reported that they wanted additional information on the edited photos that were returned, and they preferred to have a reconfirmation of whether the photo had been processed in accordance with their instructions. P5 wanted to have a binary checking result like *“privacy information cleared or not”* using the same AI

algorithm to help them confirm that they can proceed with sharing the photo. However, most participants also expressed their trust in friends and family members and mentioned that they were comfortable with it.

6.2.3 Factors Important for Deciding Whether to Share

We were also interested in what factors might be important when blind participants decide whether to share a photo. Note that because of our study setting (e.g., they already had some idea about the content of the photos they brought to the study), requesters were not making a real decision of whether to share the edited photo or not. However, participants still provided their preferences and thoughts about what factors would be important to them in this decision-making process.

Text cues: The most common factor that has been brought up by nearly all blind participants (9 out of 10) is text cues extracted from photos. Obvious text cues related to personally identifiable information and any number or ID would most likely ring a bell to blind participants and block them from sharing photos without first redacting this information with the assistance of allies. Other types of text cues would also trigger a similar reaction, including (1) extracted text that doesn't make sense (due to the imperfection of algorithms) and (2) long texts that made blind participants realize it's a scanned document that contains lots of information.

Accuracy of screening results: As discussed earlier, we discovered that although low confidence scores in AI-generated outputs can cause confusion, these scores also discourage blind participants from sharing photos and motivate them to use ImageAlly to redact private information.

Description from allies: Allies provided descriptive texts when they finished the requested task. Some blind requesters reported that they wanted to communicate with their allies again when they returned the edited photos. It was either to confirm if the allies had processed some image element specifically that allies didn't specify in the description, or to follow up the request and ask for additional assistance on the same photo. Although blind requesters could source help from the same ally using other social networking services (like directly asking them using Messages or emails), they mentioned that keeping all records on track within a single app is preferred. Unclear or unexpected descriptions from allies that require future attention would discourage our blind participants from sharing photos.

6.3 Perceived Impact on Social Relationships between Blind Individuals and Their Allies

Our third research question focuses on how the use of ImageAlly might the social relationship between blind requesters and their sighted allies. In the exit interview, we asked both user groups about how regular use of a tool like ImageAlly might

affect their social relationship. Most participants felt that ImageAlly would bring them closer to the other member of their pair. This result is in line with prior work on assistive transfer systems for solving CAPTCHAs [54], specifically, and for friendsourcing requests generally [56]. For instance, A1 noted that, as a sighted ally and friend, *"I'd have peace of mind that she is not posting anything personal."* R6 felt that ImageAlly makes it more convenient for allies to help because they can do the task remotely on their own devices rather than using blind people's devices: (*"The good thing is that it comes to them"*). R5 mentioned that ImageAlly might afford sighted people better awareness of the experiences of and challenges faced by their blind family members or friends — e.g., the limitations of image captioning systems, which blind people might rely on to make sense of images. R5 added that transferring inaccessible tasks to allies *"enhances their relationship and builds a positive connection"*. While friendsourcing requests in ImageAlly were generally considered beneficial, some participants noted that these requests should be made with good communication and respect for allies' time. For instance, R6 believed that such requests *"will be fine as long as we have good communication about doing things. I just need to be careful and not to rush them."* This result is consistent with the findings from prior work on the potential social costs of friendsourcing [14]. However, their reactions also implied that such social costs could be managed with good communication and awareness of boundaries and limits, echoing findings from prior work on friendsourcing inaccessible CAPTCHA tasks [54].

6.4 Analysis of Ally's Image Editing That Is Inconsistent with Blind Users' Preferences

Upon receiving an editing request with the blind participants' edit preferences, the sighted ally blurred parts of the picture to meet those preferences, and returned the edited picture and a description of the edits back to the blind participant. However, there could be inconsistencies between the blind participants' preferences, the ally-processed pictures, and the sighted allies' descriptions. To assess the frequency with which these inconsistencies might occur, we conducted a post-hoc analysis comparing how well ally-edited images adhered to blind users' stated preferences.

Among the 20 pictures, we found that six (four from blind participants' original pictures and two from researcher "forwarded" pictures) had inconsistencies between blind participants' preferences and how sighted allies processed them (see Table 3 in appendix). We further investigated whether the AI screening algorithms we employed would produce different results for the processed pictures than the original. This differential provided another source of data to assess whether the ally-edited images met the blind participants' preferences. Finally, we followed up with the blind participants in those six cases on how they felt about the inconsistency.

In general, the AI screening algorithms successfully cap-

tured the changes between the original and edited pictures. For example, the information on P7 and P8's cash card and ID card was mostly transcribed by optical character recognition algorithms, which also accurately reflected the remaining information in the processed pictures. The PIN number on P7's card and addresses (although detected in fragmented pieces) on P8's ID card were no longer detected in the second-time screening results. In other words, by re-screening ally-edited images, ImageAlly could help blind users confirm whether the private content they wanted to be redacted was effectively redacted by their allies.

In more carefully analyzing the six cases where we observed an inconsistency between blind users' preferences and how allies edited a picture, we observed that a key reason for these inconsistencies was that the blind participants and the ally participants had different interpretations of the former's preferences. For instance, R8's stated preference was to "blur out personally identifiable information" (PII); to that end, their ally blurred out R8's home address but not other PII (e.g., birth date, face) on the State ID card. PII can mean different things to different people, so it is possible that R8's ally had a different interpretation of PII than R8. To simplify the interface, ImageAlly currently offers only pre-defined coarse-grained options for blind users to specify their edit preferences (e.g., remove personal information). However, as we saw in this case, these coarse-grained options may leave too much open for interpretation and could cause inconsistencies between blind users' preferences and how their allies edit their images. One way to help address this issue could be encouraging blind users to provide more specific preferences.

To learn about how blind participants felt about these inconsistencies, we reached out to the blind participants R1, R5, R7, and R8 (from the table above) months after the study. We gave them the accurate descriptions of the processed pictures and asked them whether these pictures met their original expectation and if not, how they felt about this inconsistency. The accurate descriptions were generated objectively by researchers. R1 and R5 expressed that they would not mind the difference. R5 replied that "For me as long as there is some blurring on the image, that is probably fine. It shows that I am trying to protect my privacy, and usually my friends who see my postings would appreciate that." Here, R5 seemed to value more about others' impression of her attitude towards and attempt in enhancing privacy than the completeness of privacy protection.

In comparison, R7 expressed that he wished the sighted ally could "have done a better job at blurring the texts," and also gave other suggestions on sending the preferences of what to do with the pictures more efficiently. R7 said "I wish we could just call them instead of using the App to send the message (preferences in text), then they would probably know what we are talking about." Communicating nuanced preferences through text instructions could be cumbersome; voice-based communications such as phone calls may be more efficient, higher-bandwidth forms of specifying edit preferences. Therefore, complementary communication methods (e.g.,

voice calls) could be incorporated into the system to make the communication of individual privacy preferences easier.

7 Discussion

In this section, we discuss design implications from user behaviors and our human-AI system, as well as limitations and future work of this work.

7.1 Design Implications

Using ImageAlly as a design probe, our results offer a set of design implications for future private visual content management tools and assistive transfer systems for blind people.

7.1.1 Implications from User Behaviors

First, we discuss a set of implications drawn from our participants' user behaviors, including how they differed by image sources or usage scenarios, how they made sense of objects detected with low confidence scores, and how they had different privacy preferences.

User behavior may differ depending on image sources or use scenarios: Drawing on our findings where blind people behaved differently when checking their own photos versus others' photos, future designers may consider their different behaviors and design image exploring features accordingly. Specifically, they tend to use full features of image exploration when they are exploring other people's photos. In contrast, when checking on photos taken by themselves, they tend to prefer simple descriptions for efficiency.

Object detection with low confidence score might still have value: Blind people often rely much on text descriptions to inform their decisions in sending photo-editing requests and sharing photos. Drawing on our observations, when privacy-related objects are detected by algorithms albeit with low confidence scores, although it will make AI results less credible for decision-making, blind people tend to be more cautious and rely more on transferring the editing tasks, which potentially leads them into paying more attention to private content in their photos. Another related implication is that providing explanations on why the score is low might improve the user experience and give users more confidence when deciding whether to share photos or not.

Individuals can have different privacy preferences: Individuals can have different or even conflicting views on what counts as private content. A profile image might be private to some people but not to others. Moreover, views on what is private and what is not might vary across sharing contexts — for example, if one is attempting to share insurance information with a health provider, then an ID card might be an undue privacy risk. When transferring tasks, it is important to communicate preferences of what private content is and how that content should be processed. Future designs can

explore accessible ways for both parties to communicate and confirm blind users' privacy preferences.

7.1.2 Implications from ImageAlly System

Second, we discuss implications drawn from the ImageAlly prototype system itself, including discussion on usage of human-AI hybrid systems and potential customizable image double-checking for image sense-making.

Using human-AI hybrid systems: AI-based image exploration is often imperfect. Although recent advancement in computer vision has made it much easier and more robust to analyze images [32], it can still lead to confusion and thus human assistance can be useful. In comparison, purely human-based approaches to assist in visual tasks for blind people can be robust and flexible but also slow and expensive [30], which are hard to scale because of people's availability and social cost. Prior research has explored various ways of combining AI and collective human intelligence to tackle accessibility problems, such as using crowdsourcing and computer vision to detect curb ramps [21] or designing crowd-AI cameras to sense the physical world [18]. By exploring human-AI hybrid system's application to image privacy, our design probe also shed light on future designs of assistive transfer systems for blind people in managing private/sensitive visual content. Such hybrid two-layer design can be extended to many other scenarios when AI works at some level but is not perfect.

There is a spectrum of how much AI versus human work should be in this workflow. At one end of the spectrum, AI could do all the work and no humans will be involved. While perfect AI prediction is unlikely in the near future, this is theoretically possible. Previous research suggests that people who are blind tend to have a similar or higher level of privacy concern about sharing their visual content to visual question-answering systems that are powered by humans than powered by AI [45]. If the system completely relies on AI, it is presumably faster and poses less interpersonal privacy risk (since no human counterparties would see the pre-processed image), but the prediction accuracy might not be perfect. At the other end of the spectrum, only humans are involved and there could be multiple human allies involved. Dividing the image-editing task among many allies could reduce the interpersonal privacy risk of crowdsourcing because no single ally would see the entirety of the visual content. However, the task speed would reduce because the completion would depend on the schedule and work from multiple people. We view this AI vs. human design decision as trade-offs between interpersonal privacy, trust, and speed, which future research could explore further. The current, hybrid human-AI design of ImageAlly is already usable and effective, in practice, and is not contingent on any future advances in AI or computer vision.

While our blind participants mentioned in the study that ImageAlly gave them the ability to do things on their own first, the assistive transfer system approach still relies on

the interdependence between blind participants and their allies. Interdependence is considered valuable in assistive technologies [8]. ImageAlly does not necessarily change the fact that blind people might seek help from allies. Instead, it provides an integrated way for blind people to transfer the photo screening task with autonomy, and foreshadows future research on improving the accessibility of screening and editing photos for blind people when they have the needs. However, social support is not always appropriate or desired, and thus there are likely limits to its use, such as the social cost of asking for help. While our results suggest that the usage of ImageAlly could actually improve the social relationship between blind requesters and sighted allies, future research could further examine the cost of social support in such systems.

Consider customizable image double-check: Some blind participants wanted to be able to check the edited photos after receiving them back from allies. Specifically, there was a desire to compare the AI screening results before and after their allies' editing. This suggests the option of double checking the ally edited photos and highlighting differences before and after human editing. For instance, AI can be applied again to the ally edited photos and can simply say for instance "the human faces are no longer present." Furthermore, another implication we drew from participants' responses is that there is value in making the double check process customizable. For example, requesters can set a rule like "Face on the right side should not be detectable" and thus, both allies and the algorithm would have a clear metric of what to detect and edit.

7.2 Limitations and Future Work

7.2.1 Limitation of The Lab Study

ImageAlly has some limitations when deployed in our lab study. For example, we found that requesters felt the system did not provide sufficient notifications to them when allies receive the message and start working on it. Participants reported that they prefer to have a way to know if their allies started processing the photos so that they have a better sense of whether to send requests to another ally. Another limitation was that requesters need to select a contact to send the request. Participants wanted more flexible ways to select one or more contacts when they needed to, like maintaining a commonly used friend list. We also did not have a large sample size. It is challenging to recruit blind participants, and it was more so in our study because every session requires a pair of a blind person and an ally. However, our sample size is on par with other privacy/security user studies focusing on blind people [7, 24, 54].

Another limitation stems from the controlled nature of our lab study. We chose to conduct an exploratory lab study rather than a field deployment because we were at an exploratory stage of designing such hybrid human-AI system and we want to obtain rich and qualitative data on users' perceptions of and reactions to ImageAlly as a probe. Also, field deployment is

more appropriate in a later stage of the iterative design process. In the meantime, we also recognize that some study settings like asking participants to imagine how ImageAlly would affect their relationship are limited and can only be answered in a lab study but in a field deployment. We consider this as a promising future work. As a result, the images used in our lab study are limited in representing the types of photos blind people might share in real lives. Additionally, although we told allies that requests from their blind friends or family members may come at any time, and that they could do whatever they pleased in the meanwhile rather than waiting for ImageAlly requests, allies still put themselves in a lab study situation and were always available when requesters sent requests. However, in practice, allies might not always be available at the time when blind people need help. However, compared to prior work on transferring CAPTCHA tasks (which usually expire in 2 minutes) [54], photo screening is often less urgent and thus is an asynchronous task, which can also be sent to more than one ally at the same time. This could help scale ImageAlly since ally's availability is less of a concern.

7.2.2 Limitations of The ImageAlly Prototype

ImageAlly was implemented as a proof-of-concept design probe rather than as a full-fledged production-ready system, and its current implementation is not bulletproof for privacy and security. While blind participants were all comfortable with choosing a trusted ally to deal with their photos, the system might be exploited by malicious attackers. For example, since requests were sent to allies using URLs via SMS text messages (it will come from a phone number used by ImageAlly), it might be intercepted by malicious third parties. In addition, attackers might send malicious requests to unsuspecting allies and get them to edit photos for free. Since ImageAlly was built as a friendsourcing system, requesters and allies should already have a trustworthy relationship. Several strategies could help mitigate the above privacy/security risks, including, for example, requiring registration and authentication on both sides, building a trustable contact list (whitelist), setting request quotas per day, and each party sending a separate confirmation text message to the other party directly using their own phone number. Apart from the security risks of using ImageAlly's transferring feature, there are also privacy and security implications of using 3rd party APIs. This risk could be mitigated by avoiding using 3rd party commercial APIs and developing proprietary machine learning models based on datasets like VizWiz. These are implementation-specific trade-offs, and not fundamental risks imposed by the system design.

Another limitation came from the fact that ImageAlly was built on existing computer vision models to detect objects in visual content. Although our intention was to provide users with full agency and control of their own visual data by listing all possible objects detected to empower their image editing/sharing decision-making, existing object detection

algorithms can have false negative errors (e.g., a card on the table was not detected because it was far away from the camera). We consider such inaccuracies as motivations for future AI solutions. Also, image analysis contains a variety of means beyond what we proposed in section 4.1, we consider studying what image analysis is necessary and efficient for blind users to make sense of pictures as new challenges for future work.

7.2.3 Future Work

There are several areas of future work that this work opens. One potential avenue is the exploration of the critical role that allies play in our human-AI system. While our current research scope focuses on blind users, it is necessary to study the overall satisfaction of allies and how they are impacted with unintended social tensions. We can also explore how such system can encourage allies to engage in visual tasks related to blind individuals' privacy. Another potential direction is to use more advanced computer vision techniques to better study the dynamics between human and machine intelligence with a more robust and reliable system. Furthermore, future research can also study additional use cases beyond sharing on social media, as ImageAlly could be used in virtually any case where blind users want to send or share visual content with another party (e.g., uploading an image for registration or reimbursement).

8 Conclusion

To assist blind people in detecting and redacting private content in photos that they might consider sharing online, we designed and implemented a proof-of-concept probe — ImageAlly. ImageAlly employs a hybrid human-AI workflow that affords blind users AI-generated insights about potential private content in images, and then facilitates the solicitation of targeted editing assistance from trusted allies. Through an exploratory lab study with recruited pairs of blind people and their sighted allies, we found that both parties liked ImageAlly. We also found that blind users preferred coarse, minimal descriptors for their own photos (e.g., number of faces detected) but more fine-grained descriptors on others' photos from the AI-generated screening results. Furthermore, our results suggest that use of assistive transfer systems like ImageAlly has the potential to strengthen the relationship between blind requesters and their sighted allies. ImageAlly could also increase sighted allies' awareness of the challenges faced by their blind friends and family members.

9 Acknowledgements

We thank our participants for their contributions and sharing their insights. This research was in part supported by the National Science Foundation (NSF) grants #2126314, #2028387, #2125925, and #2148080 and the CRA CIFellows program.

References

- [1] Tousif Ahmed, Roberto Hoyle, Kay Connelly, David Crandall, and Apu Kapadia. Privacy Concerns and Behaviors of People with Visual Impairments. In *CHI2015*, pages 3523–3532, 2015.
- [2] Microsoft Seeing AI. <https://apps.apple.com/us/app/seeing-ai/id999062298>, Accessed: 2022-01-31.
- [3] Aira. <https://aira.io/>, Accessed: 2022-01-31.
- [4] 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.
- [5] Jyoti Aneja, Aditya Deshpande, and Alexander G Schwing. Convolutional image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5561–5570, 2018.
- [6] Giulia Barbareschi, Catherine Holloway, Katherine Arnold, Grace Magomere, Wycliffe Ambeyi Wetende, Gabriel Ngare, and Joyce Olenja. The social network: How people with visual impairment use mobile phones in kibera, kenya. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2020.
- [7] Nata Barbosa, Jordan Hayes, and Yang Wang. UniPass: Design and Evaluation of A Smart Device-Based Password Manager for Visually Impaired Users. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp 2016)*, 2016.
- [8] 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.
- [9] Cynthia L Bennett, Jane E, Martez E Mott, Edward Cutrell, and Meredith Ringel Morris. How teens with visual impairments take, edit, and share photos on social media. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, pages 1–12, 2018.
- [10] 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.
- [11] Rupert RA Bourne, Jaimie Adelson, Seth Flaxman, Paul Briant, Michele Bottone, Theo Vos, Kovin Naidoo, Tasanee Braithwaite, Maria Cicinelli, Jost Jonas, et al. Global prevalence of blindness and distance and near vision impairment in 2020: progress towards the vision 2020 targets and what the future holds. *Investigative Ophthalmology & Visual Science*, 61(7):2317–2317, 2020.
- [12] Richard E Boyatzis. *Transforming qualitative information: Thematic analysis and code development*. sage, 1998.
- [13] Michael Boyle, Christopher Edwards, and Saul Greenberg. The effects of filtered video on awareness and privacy. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work*, pages 1–10, 2000.
- [14] Erin L Brady, Yu Zhong, Meredith Ringel Morris, and Jeffrey P Bigham. Investigating the appropriateness of social network question asking as a resource for blind users. In *Proceedings of the 2013 conference on Computer supported cooperative work*, pages 1225–1236, 2013.
- [15] John Brooke. SUS-A quick and dirty usability scale. *Usability evaluation in industry*, 189(194):4–7, 1996.
- [16] Lookout by google. https://play.google.com/store/apps/details?id=com.google.android.apps.accessibility.reveal&hl=en_US&gl=US, Accessed: 2022-01-31.
- [17] Arianna Dagliati, Simone Marini, Lucia Sacchi, Giulia Cogni, Marsida Teliti, Valentina Tibollo, Pasquale De Cata, Luca Chiovato, and Riccardo Bellazzi. Machine learning methods to predict diabetes complications. *Journal of diabetes science and technology*, 12(2):295–302, 2018.
- [18] Anhong Guo, Anuraag Jain, Shomiron Ghose, Gierad Laput, Chris Harrison, and Jeffrey P Bigham. Crowd-ai camera sensing in the real world. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(3):1–20, 2018.
- [19] 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.
- [20] Danna Gurari, Yanan Zhao, Meng Zhang, and Nilavra Bhattacharya. Captioning images taken by people who are blind. In *European Conference on Computer Vision*, pages 417–434. Springer, 2020.

- [21] Kotaro Hara, Jin Sun, Robert Moore, David Jacobs, and Jon Froehlich. Tohme: detecting curb ramps in google street view using crowdsourcing, computer vision, and machine learning. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*, pages 189–204, 2014.
- [22] Susumu Harada, Daisuke Sato, Dustin W Adams, Sri Kurniawan, Hironobu Takagi, and Chieko Asakawa. Accessible photo album: enhancing the photo sharing experience for people with visual impairment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2127–2136, 2013.
- [23] Rakibul Hasan, Eman Hassan, Yifang Li, Kelly Caine, David J Crandall, Roberto Hoyle, and Apu Kapadia. Viewer experience of obscuring scene elements in photos to enhance privacy. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2018.
- [24] 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.
- [25] Simao Herdade, Armin Kappeler, Kofi Boakye, and Joao Soares. Image captioning: Transforming objects into words. *arXiv preprint arXiv:1906.05963*, 2019.
- [26] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, et al. Technology probes: inspiring design for and with families. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 17–24, 2003.
- [27] Panagiotis Ilia, Iasonas Polakis, Elias Athanasopoulos, Federico Maggi, and Sotiris Ioannidis. Face/off: Preventing privacy leakage from photos in social networks. In *Proceedings of the 22nd ACM SIGSAC Conference on computer and communications security*, pages 781–792, 2015.
- [28] Masakazu Iwamura, Naoki Hirabayashi, Zheng Cheng, Kazunori Minatani, and Koichi Kise. Visphoto: Photography for people with visual impairment as post-production of omni-directional camera image. page 1–9. ACM, Apr 2020.
- [29] Chandrika Jayant, Hanjie Ji, Samuel White, and Jeffrey P Bigham. Supporting blind photography. In *The proceedings of the 13th international ACM SIGACCESS conference on Computers and accessibility*, pages 203–210, 2011.
- [30] Ece Kamar. Directions in hybrid intelligence: Complementing ai systems with human intelligence. In *IJCAI*, pages 4070–4073, 2016.
- [31] Ioannis Kavakiotis, Olga Tsave, Athanasios Salifoglou, Nicos Maglaveras, Ioannis Vlahavas, and Ioanna Chouvarda. Machine learning and data mining methods in diabetes research. *Computational and structural biotechnology journal*, 15:104–116, 2017.
- [32] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything, 2023.
- [33] Pavel Korshunov, Andrea Melle, Jean-Luc Dugelay, and Touradj Ebrahimi. Framework for objective evaluation of privacy filters. In *Applications of Digital Image Processing XXXVI*, volume 8856, pages 265–276. SPIE, 2013.
- [34] Yifang Li, Nishant Vishwamitra, Hongxin Hu, Bart P Knijnenburg, and Kelly Caine. Effectiveness and users’ experience of face blurring as a privacy protection for sharing photos via online social networks. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 61, pages 803–807. SAGE Publications Sage CA: Los Angeles, CA, 2017.
- [35] Yifang Li, Nishant Vishwamitra, Bart P Knijnenburg, Hongxin Hu, and Kelly Caine. Effectiveness and users’ experience of obfuscation as a privacy-enhancing technology for sharing photos. *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW):1–24, 2017.
- [36] Yiyi Li and Ying Xie. Is a picture worth a thousand words? an empirical study of image content and social media engagement. *Journal of Marketing Research*, 57(1):1–19, 2020.
- [37] Haley MacLeod, Cynthia L Bennett, Meredith Ringel Morris, and Edward Cutrell. Understanding blind people’s experiences with computer-generated captions of social media images. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 5988–5999, 2017.
- [38] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for csw and hci practice. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–23, 2019.
- [39] Richard McPherson, Reza Shokri, and Vitaly Shmatikov. Defeating image obfuscation with deep learning. *arXiv preprint arXiv:1609.00408*, 2016.

- [40] Deirdre K Mulligan and Helen Nissenbaum. The concept of handoff as a model for ethical analysis and design. In *The Oxford Handbook of Ethics of AI*, page 233. Oxford University Press, 2020.
- [41] Be my eyes - see the world together. <https://www.bemyeyes.com/>. Accessed: 2022-01-31.
- [42] Pedro J Navarro, Carlos Fernandez, Raul Borraz, and Diego Alonso. A machine learning approach to pedestrian detection for autonomous vehicles using high-definition 3d range data. *Sensors*, 17(1):18, 2017.
- [43] Cognitive services — APIs for AI developers. <https://azure.microsoft.com/en-us/services/cognitive-services/>, Accessed: 2022-01-31.
- [44] Rachel N Simons, Danna Gurari, and Kenneth R Fleischmann. " i hope this is helpful" understanding crowdworkers' challenges and motivations for an image description task. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2):1–26, 2020.
- [45] Abigale Stangl, Kristina Shiroma, Bo Xie, Kenneth R Fleischmann, and Danna Gurari. Visual content considered private by people who are blind. In *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, pages 1–12, 2020.
- [46] Testflight. <https://developer.apple.com/testflight/>, Accessed: 2022-01-31.
- [47] Marynel Vázquez and Aaron Steinfeld. Helping visually impaired users properly aim a camera. In *Proceedings of the 14th international ACM SIGACCESS conference on Computers and accessibility*, pages 95–102, 2012.
- [48] Nishant Vishwamitra, Bart Knijnenburg, Hongxin Hu, Yifang P Kelly Caine, et al. Blur vs. block: Investigating the effectiveness of privacy-enhancing obfuscation for images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 39–47, 2017.
- [49] Accessibility - Vision. <https://www.apple.com/accessibility/vision/>, Accessed: 2022-01-31.
- [50] Violeta Voykinska, Shiri Azenkot, Shaomei Wu, and Gilly Leshed. How blind people interact with visual content on social networking services. In *Proceedings of the 19th acm conference on computer-supported cooperative work & social computing*, pages 1584–1595, 2016.
- [51] 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.
- [52] Shaomei Wu, Jeffrey Wieland, Omid Farivar, and Julie Schiller. Automatic alt-text: Computer-generated image descriptions for blind users on a social network service. page 1180–1192. ACM, Feb 2017.
- [53] Zhongliang Yang, Yu-Jin Zhang, Sadaqat ur Rehman, and Yongfeng Huang. Image captioning with object detection and localization. In *International Conference on Image and Graphics*, pages 109–118. Springer, 2017.
- [54] Zhuohao Zhang, Zhilin Zhang, Haolin Yuan, Natã M Barbosa, Sauvik Das, and Yang Wang. Webally: Making visual task-based captchas transferable for people with visual impairments. In *Seventeenth Symposium on Usable Privacy and Security (SOUPS) 2021*, pages 281–298, 2021.
- [55] Yuhang Zhao, Shaomei Wu, Lindsay Reynolds, and Shiri Azenkot. The effect of computer-generated descriptions on photo-sharing experiences of people with visual impairments. *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW):1–22, 2017.
- [56] 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.

A Appendix

A.1 Non-Visual Interface Details

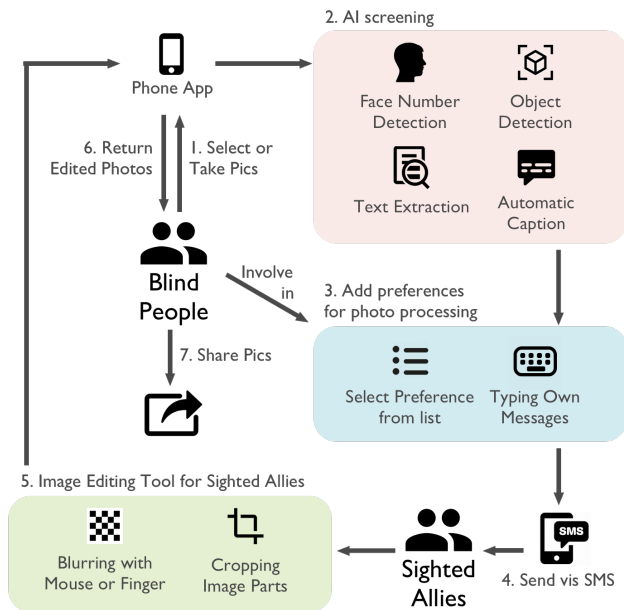


Figure 2: Workflow of ImageAlly: Blind users select or take a picture from the App, uses AI to screen the photo, and read through the results. Then they will be prompted to add a preference of what to do with the photo and choose whether to send it via SMS messages to sighted allies. If they choose to send a request, allies will use an interactive image editing tool to blur or crop out image parts and send it back with descriptions of what they did. Blind people will then have the option to share it with others or on social media.

A.1.1 Detection of the Quantity of Faces:

Whether the number of faces matches blind people’s own expectation is an important measurement of whether the image has unnecessary private or sensitive information. For example, if a blind user is trying to take a selfie, a family photo, or photos of scenery, they may have specific expectations of how many faces should appear in the photo. If the number of detected faces diverges from this expectation, the blind user may elect for further screening and processing of the photo.

A.1.2 Detection of Related Objects:

State-of-the-art computer vision models can identify objects in images and describe these objects in natural language. We leverage commercial APIs [43] of state-of-the-art models to provide both object category names and detection confidence

scores to help blind users decide whether to transfer the images to human allies for additional processing. For example, if an ID card is detected in the photo, ImageAlly will present the user with a prompt akin to the following: “We have detected the following objects in the picture, together with a percentage number showing how confident we are for each detected object: Text (72% sure), Card (81% sure).” Note that the categories are provided by the Microsoft Azure Object Detection API [43] and therefore we are unable to get a full category list of objects to be detected. We consider this out of our scope because our ImageAlly design is intended to be able to generalize for use with other current and future improved object detection models.

A.1.3 Detection of Related Texts:

If text is detected, we then employ a commercial optical character recognition (OCR) API [43] to extract the text and present it to the user.

A.1.4 Automatic Captioning:

We also use commercial neural image captioning APIs to generate a caption [43], along with an associated confidence score, to help users generally and broadly understand the broad strokes of what is captured in the image. An example caption: “Additionally, we are 93.27% sure that this picture can be generally described as: graphical user interface, text, application, chat or text message.”

A.1.5 Detection of Adult Content:

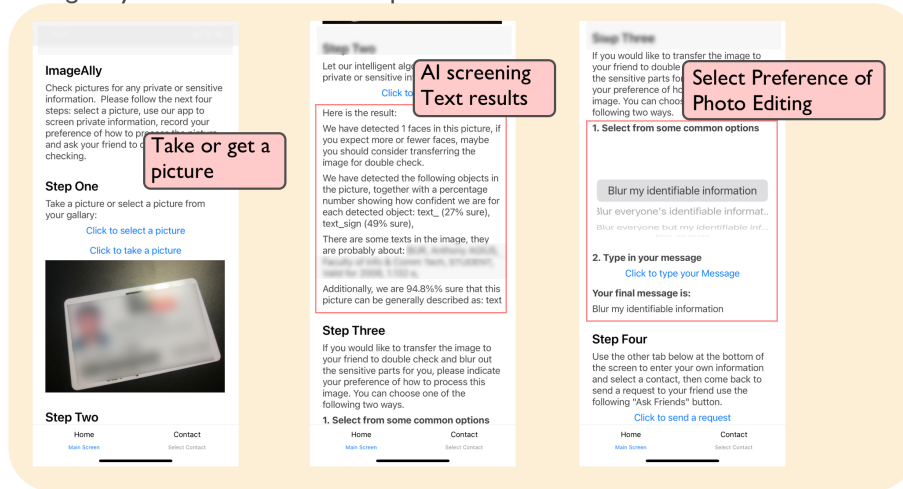
Using Microsoft’s visual feature APIs [43], we also added adult content detection.

Note that we are providing as many image analysis result as possible for screening processes. With the development of image analysis techniques, this subsection can go longer as needed. We are using these detection categories to provide as an example in designing ImageAlly probe. Furthermore, the use of AI in ImageAlly is only to provide descriptions of *potentially* private content. It is ultimately up to the blind user to determine if they deem content to be private and so if it should be edited or left alone.

A.2 Pilot Study

Prior to the main study, we conducted a series of pilot sessions with seven blind participants to improve the workflow and accessibility of ImageAlly. We conducted these pilots similar to our full study with our ImageAlly prototype. Specifically, we did four pilot study sessions with seven participants using a task-based usability test.

ImageAlly Interface for Blind People



Interface for Allies

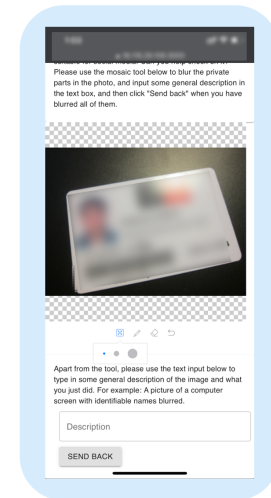


Figure 3: Interfaces of ImageAlly. Left side, UI for blind users: (1) select or take an image, (2) Using images selected in Step one, screen the image using AI algorithms (Note that the detected result of the ID card in Step one is considered as “text” and “text-sign” (hand-written text sign) because the ID card mostly contained text. The detection result is limited to existing commercial APIs, which will be discussed in section 7.2.2), and (3) set image editing preferences and then send a request to allies (The common preference options listed here include blurring out the requester’s or everyone’s identifiable information or faces. Requesters can also indicate their own preferences). (b) UI for allies: Description of the task, together with AI screening results and a text input for describing their actions.

A.2.1 Pilot Study Method

The first session included one blind participant and one of our researchers acted as the ally upon the participant’s consent. The remaining three pilot sessions included one blind participant and one of their sighted friends or family members as an ally. We used our initial ImageAlly prototype as the apparatus for the study, and prepared a image for the participants to use as the material. The image was a mobile screenshot of a work group chat, containing private information like co-workers’ names, work content, and their profile avatar images. During each study session, the participants were divided into two Zoom breakout rooms so that they could focus on testing ImageAlly without talking to each other and causing disturbance. First the researchers introduced the study and the task, followed by a series of interview questions about their previous experience of receiving or providing visual assistance regarding photos. Then researchers asked the participants to use ImageAlly to process the prepared photo before they share it to a third party. As the procedure of using ImageAlly in pilot study is identical to the procedure of formal study evaluation, we present in Section 5.3. We recorded and transcribed the data for each session. Two researchers developed themes from the transcripts using thematic analysis. Then researchers met online to discuss how participants’ feedback informed improvement of ImageAlly and iterated the system accordingly.

A.2.2 Pilot Study Results

Below we summarize what we learned from the pilot study and the main changes we made accordingly to the system.

Improve interface accessibility. Pilot study participants reported that ImageAlly’s interface could be made more accessible. For example, participants mentioned that the navigation inside the ImageAlly probe could be improved with heading and page-based navigation. They also suggested using prompts and confirmations more often when they or their allies finish a certain step. Specifically, they wanted to receive notifications when their allies received the request and started working on it. Based on this feedback, we added heading navigation to the prototype and added notifications using accessible pop-up alerts each time a user completes a step.

Present AI results to allies. Participants also suggested that their allies should view the same screening results from AI as they themselves did. By presenting the AI results to allies, they could understand what their blind friends are getting and can then make better decisions when they edit the photos. For example, there might be privacy-related content in the photo that was not detected by AI, or inaccurate AI predicted results that blind users obtain and use to make decisions. Therefore, we present the AI results to allies when they receive the image processing request.

Improve asynchronous communications. In our original prototype, the communication happened in an unidirectional

way. Only blind people were able to articulate preferences for what they like their ally to do with their photos. However, participants suggested that allies should also be able to respond with what they did to the photos. Based on this, we added a bidirectional communication channel to allow allies to describe the edits they made to blind requesters.

A.3 Participants Biographics

We provide participants' biographics here in table 1.

A.4 Photos Used by Participants

We provide general descriptions of photos (table 2) that contained (outdated) private information of blind participants.

A.5 Comparison Between Inconsistent Image Editing

We provide a comprehensive analysis result of comparison between inconsistent image editing results transferred back from sighted allies here in table 3.

A.6 Pre-study Questions

A.6.1 Questions for Blind Requesters

- Do you take photos? When and how? What kind of photos?
- Do you share your photos with others? When and how? What kind of photos?
- Do you share your photos on social media? When and how? What kind of photos?
- Do you edit your photos before you share? What do you edit photos for and how do you do that?
- How do you decide what photos to share?
- What's usually in the picture parts where you want to edit? How do you usually do that?

A.6.2 Questions for Sighted Allies

- Did your friend/family member (blind or low vision) ever consult you about photos they take? Like whether the image contains the right content, whether the figure looks good, whether there's private information that should be cropped out or blurred out?
- Do you have any concerns seeing their private information in the photos if there's any?
- What's your preference for how to be contacted by the requester? Like text messages, email, DMs from social media App, etc.

A.7 Exit Interview Questions

A.7.1 Questions for Blind Requesters

- Please give us a general impression of the idea and process
- To recap the screening results, how do you interpret the AI outputs? What's useful? What's not useful?
- What do you expect to see more in the AI result for better decision-making on whether to share it to public?
- Have you tried other tools that help you recognize the image contents? What kind of information helps you decide whether there is private or sensitive information?
- How do you make use of the preference recording function (step3)?
- How do you make use of the returned image from friends? Do they meet your expectations?
- Would using this app influence your relationship with [the other party]? How would it potentially affect the relationship in any positive or negative way?
- Do you have any more suggestions?
- Please rate from 1-5 (strongly disagree to strongly agree) for the following statements

- I think that I would like to use this system frequently.
- I found the system unnecessarily complex.
- I thought the system was easy to use.
- I think that I would need the support of a technical person to be able to use this system.
- I found the various functions in this system were well integrated.
- I thought there was too much inconsistency in this system.
- I would imagine that most people would learn to use this system very quickly.
- I found the system very cumbersome to use.
- I felt very confident using the system.

A.7.2 Questions for Sighted Allies

- Please give us a general impression of the idea and process
- How do you think about the instruction by the requester? Is it helpful?
- How do you think about the image editing function?
- How do you think about using SMS text to transfer the request?

Table 1: Blind participants' demographics, including their age group, gender identity, self-described disability, their allies' gender and relationships (allies' relationship to requesters).

Requester	Age	Gender	Self-Described Disability	Ally	Ally Age	Ally Gender	Relationship
R1	35-44	Female	Blind	A1	35-44	Male	Friend
R2	25-34	Female	Blind	A2	55-64	Female	Mother
R3	25-34	Female	Blind	A3	25-34	Male	Friend
R4	18-24	Male	Blind and Hearing Impairments	A4	18-24	Female	Partner
R5	45-54	Female	Blind	A5	18-24	Female	Daughter
R6	55-64	Female	Blind and Hearing Impairments	A6	55-64	Female	Friend
R7	55-64	Male	Blind	A7	55-64	Female	Sister
R8	35-44	Male	Blind	A8	35-44	Female	Friend
R9	25-34	Male	Blind	A9	25-34	Female	Friend
R10	18-24	Male	Blind	A10	18-24	Male	Brother

Table 2: Photos created and used by blind participants

Participant ID	Photo they chose
R1	A man in front of a birthday cake
R2	Mobile phone screenshot with server port and password
R3	Mobile phone screenshot with calendar invite and personal information
R4	Room surroundings
R5	Medical bottle with prescriptions
R6	Insurance document on the table
R7	Expired cash card on the table
R8	Expired state ID card on the table
R9	Transaction screenshot with transaction ID and part of bank account number
R10	ID card on the table

- How do you make use of the text input as a way to inform requesters about what you edited?
 - I think that I would need the support of a technical person to be able to use this system.
- Do you have any more suggestions?
 - I found the various functions in this system were well integrated.
- Please rate from 1-5 (strongly disagree to strongly agree) for the following statements
 - I think that I would like to use this system frequently.
 - I found the system unnecessarily complex.
 - I thought the system was easy to use.
- I thought there was too much inconsistency in this system.
- I would imagine that most people would learn to use this system very quickly.
- I found the system very cumbersome to use.
- I felt very confident using the system.

PID	Which Picture	Criteria	Description or Quote
R1	Original	Picture Abstract Blind User Preference Processed Image Ally's description AI Screening Difference	A man smiling in front of a birthday cake with candles on it "Blur out the cake" The candles on the cake were blurred "I blurred out the cake" Candle is no longer detected in the processed image, but cake still is
R5	Original	Picture Abstract Blind User Preference Processed Image Ally's description AI Screening Difference	A medicine bottle with prescriptions on it, including name, tablet size, prescription, ID, and date "Blur out prescription information" The name, tablet size, ID and date were blurred, but the prescription including the medical condition was not blurred "I blurred out the prescription" Texts are no longer detected in the processed image
R7	Original	Picture Abstract Blind User Preference Processed Image Ally's description AI Screening Difference	A picture of a cash card's back, including security information of card number and PIN "Blur out security information of this card" The PIN was blurred out, but the card number is not "I erased the password" Text changed from instructions and card number and password to instructions and card number only
R8	Original	Picture Abstract Blind User Preference Processed Image Ally's description AI Screening Difference	A state ID card's front page, including name, date of birth, address, biometrics, and face picture "Blur out personal identifiable information" The address was blurred out, but the rest of the information was not "I blurred your address" Text of address is no longer detected in the processed image
R4	Forwarded	Picture Abstract Blind User Preference Processed Image Ally's description AI Screening Difference	Mobile screenshot of a work group chat history with co-workers' names and profile pictures "Blur out personal identifiable information" Only the names were blurred, the profile pictures still remain in sight "I blurred their information" Text of names were no longer detected
R6	Forwarded	Picture Abstract Blind User Preference Processed Image Ally's description AI Screening Difference	Mobile screenshot of a work group chat history with co-workers' names and profile pictures "Blur out colleagues' info" Only the profile pictures were blurred, the names still remain in sight "I blurred out their faces" Face number changed to 0

Table 3: Comparison between the pictures that contain inconsistent editing with ally processed pictures, including (1) what this picture was about, (2) what was the blind participant's preference for processing the picture, (3) what was the processed picture like, (4) what was the described by allies, and (5) what were the differences between AI screening results for the two pictures.