Display Leakage and Transparent Wearable Displays: Investigation of Risk, Root Causes, and Defenses

Tadayoshi Kohno
University of Washington

Joel Kollin
Microsoft

David Molnar
Microsoft

Franziska Roesner
University of Washington

Abstract

Transparent near-eye displays are shipping now for augmented reality applications. In addition to these applications, they promise a private display safe from shoulder surfing. Multiple researchers in the security and HCI communities have proposed systems building on the assumption these displays are private [14, 23, 24]. Unfortunately, this assumption is not always true. We find multiple shipping displays suffer from *display leakage*: an adversary who observes a user wearing the display can reconstruct the contents of the display from light leaked by the "outward-facing" part of the display. We propose defenses against display leakage and analyze them in context of a range of display designs.

1 Introduction

Transparent (see-through) displays are a core element of emerging, wearable computing devices. These displays allow the user of the device to simultaneously see information on the display *and* see the world beyond. One of the most well-known examples of this technology is the Google Glass. Figure 1 shows an image of a user wearing the Glass. Numerous other transparent displays exist, including (for example), the Silicon Micro Display ST-1080, the Epson Moverio BT 200, the Lumus DK-40, and the Meta One.

These wearable, near-eye displays promise a user safety from shoulder surfing attacks, in which a bystander observes the content of a user's screen. Multiple sets of researchers have proposed applications that take advantage of such private wearable displays (e.g., for password managers that display passwords [14], applications that overlay the real world with sensitive, decrypted content [16], and smartphone PIN entry [23]). All of these applications assume that the contents of the display cannot be observed by bystanders. Unfortunately, this assumption is not always true. We show that multiple classes of displays suffer from *display leakage*, in which light from the "outward-facing" part of the display can be analyzed to reveal information about the contents of the display.

Figure 1: On top, an active Google Glass. On bottom, a recent proposal for smartphone unlocking that assumes the Glass is a private display [23].

At present, wearable display manufacturers are focused on improving transparent wearable displays along traditional metrics such as size, cost, power consumption, resolution, and contrast [8]. We propose another metric, which is not included in Kress and Starner's list of goals [8]: *display privacy*. We say that a device has poor display privacy if it leaks information about the content that a user is viewing to someone on the other side of the display; this exposure of visual information is *display leakage*.

Consider Figure 1. On the top, a picture of the Google Glass worn by a person with the display active. Observe the visible glow in the user's display. On the bottom, an illustration of a recent proposal for smartphone unlocking that assumes the Glass is a private display and uses the Glass to display information only the user can see. Of course, the Glass is also used for viewing email, visiting web sites, and other day to day activities.

The glow in the Glass's transparent display suggests at least a portion of what the user is viewing may be visible to third parties. A key question: what and how much information is available, and how concerned should users be? Absent a rigorous study, although it is possible to speculate, it is impossible to know for sure. We fill this gap and experimentally study these questions in this paper, using two example technologies as our reference points: the Google Glass and the Silicon Micro Display ST-1080.

Our experiments suggest that both technologies do, indeed, leak what we conclude to be a concerning amount of information to third parties on the other sides of the displays, including at a minimum information about what activities a user might be doing. We have put all the raw images that we collected during this study on `http://www.displayleak.com`, for reviewers to evaluate independently; for this submission, we have removed or cropped images that might de-anonymize the authors. We also show that two other technologies — the Meta One and a Lumus eyepiece display — suffer from display leakage, though we do not study them in depth.

We then proceed to study the underlying physical properties of these devices that lead to display leakage. While understanding the physical inner-workings of these devices is not new knowledge — the device designers certainly knew how the displays were designed — our perspective is different since (1) we evaluate the physical properties that lead to display leakage, not that lead to high-quality displays for the users (the latter being our assessment of the primary focus of prior works) and (2) we are (to our knowledge) the first to share this information with the computer security community. Moreover, (3) we propose a simple defense against the attack using the properties of polarized light, and we validate this defense with a real transparent head mounted display. We also view this line of inquiry as core to the emerging technologies computer security research community, and complementary to efforts focused on improving the security and privacy of future wearable devices and their applications, such as PlaceAvoider [18] and other related works [6, 7, 15], as well as essential to understanding the risks of works that use transparent wearable displays to render private content [14, 16].

A naive solution to protecting display privacy might be to make the displays opaque, but doing so would compromise one of the key functionality goals for transparent displays. By leveraging our cross-disciplinary backgrounds in computer security and optical engineering, we find that we are able to provide both display privacy and display transparency by altering the optical pathways within the devices.

Our main approach builds on the properties of polarized light. At the highest level, our defense polarizes the light entering the display in one direction, then add another polarizer in the opposite direction on the outside of the transparent display. This means that light leaving the transparent display cannot be reconstructed into a coherent image. While implementing this solution does require changing the optical pathway, we argue that the required changes can be cheap — adding an absorbing polarizer to a wearable device costs on the order of a few dollars — and even aftermarket. The approach we tested does make the outside world seem dimmer when wearing the display, similar to wearing polarized sunglasses. However, we argue that this is an acceptable tradeoff for providing a private display, especially for form factors such as Google Glass that do not cover the user's entire field of view. Alternatively, the polarizing filter could be used only selectively when the user is in non-private environments or viewing sensitive content, similarly to how laptop privacy filters are used today.

We implemented our polarization defense with the Silicon Micro Display, validating that our approach in fact reduces display leakage. We also sketch an alternative approach, which uses narrowband illumination of an LCoS microdisplay and narrowband filters to prevent the display light from being seen by a third party.

Our contributions are the following:

1. We formulate *display privacy* as a key goal for transparent wearable displays.
2. We experimentally analyze the degree to which two example transparent wearable displays violate the display privacy goal, i.e., we experimentally evaluate their *display leakage* characteristics.
3. We reconstruct the physical and optical properties of these displays and, in doing so, determine why they leak display information to third parties on the other sides of the displays.
4. We develop and validate a simple, low-cost defense that results in both transparent and private displays; our defense changes the optical pathways in the devices.

In Section 2 we provide background on transparent head-mounted displays and augmented reality, as well as information leakage via the visual channel. Then in Section 3 we describe our experiments with measuring display leakage for two shipping head-mounted displays, the Google Glass and the Silicon Micro Display ST-1080. After we establish that these are vulnerable to display leakage, in Section 4 we review the optics behind these designs, as well as discuss alternative optic designs available on the market today. In Section 5 we describe modest hardware changes that can reduce display leakage, and we validate our approach. Finally, we conclude in Section 6.

2 Background

Augmented Reality and Wearable Displays. Sutherland described a head-mounted display showing three-dimensional information in 1968, together with mechanical and ultrasonic methods for determining head position in order to update the

image [17]. Since then, decades of work by hundreds of researchers have focused on different aspects of head-mounted displays, as well as applications for these displays, as surveyed by Azuma [1]. Transparent head-mounted displays are an important piece of the puzzle because they enable overlaying virtual content on top of real world objects. Display leakage, however, has not been a primary concern; instead work has focused on other matters, such as field of view.

Applications Assuming Private Wearable Displays. A head-mounted display promises the additional benefit that what is shown to the user cannot be seen by others. Ofek et al. investigate this in the context of prompting people with auxiliary information during a conversation, under the assumption that the prompt cannot be seen by the conversation partner [12]. Simkin et al. propose the use of wearable displays to enable real-world applications of cryptography, such as encrypted paper documents that are decrypted only in the wearable display [16]. Roesner et al. propose an "augmented password manager" that recognizes the website seen by the user, then displays the user's password in a transparent display [14]. Winkler et al. [23] and Yadav et al. [24] propose unlocking mechanisms that assume the Google Glass is a private display. These are part of a growing body of work that proposes novel security mechanisms assuming what is shown on a near-eye display is not viewable by others near the user. Our work shows that *this assumption is false* for several classes of displays shipping today.

Information Leakage via the Visual Channel. Visual eavesdropping can leak sensitive information. The term "shoulder surfing" refers to the practice of looking over a user's shoulder while she views sensitive content [22]. Shoulder surfing attacks are often targeted at obtaining passwords or PINs, and recent research has demonstrated the utility of wearable devices like the Google Glass in shoulder surfing attacks [5]. Companies such as 3M sell "privacy filters," polarizing screens that reduce the effective angle at which a monitor can be viewed [4], explicitly to limit shoulder surfing. Head-mounted displays offer the promise of eliminating shoulder surfing by providing a private display unique to the user, but display leakage undermines this promise.

Even if a monitor is facing away from an adversary, stray reflections can unintentionally reveal sensitive information. Backes et al. investigated the recovery of information from reflections off eyes, teapots, and other reflective surfaces [2, 3]. Our work, in contrast, focuses on the characteristics of the display itself. In our setting, no reflections off incidental surfaces are needed, because the display leaks light outwards toward an adversary. While fixing this display leakage does not make reflections off eyeballs go away, our work shows that common transparent head-mounted display designs leak substantial information without even requiring sophisticated image processing.

Figure 2: The Silicon Micro Display we use in action. Note the reflection from the lenses at this angle.

3 Analysis of Display Leakage

We analyzed the display leakage properties of two archetypical examples of commercially available transparent displays: the Google Glass and the Silicon Micro Display. Figure 1 shows a picture of the Glass, and Figure 2 shows a picture of the Silicon Micro Display ST-1080. The Glass is a wearable (head-mounted) device with a small transparent display placed above the right eye. The Glass is wireless and controlled primarily by voice inputs and a small touch pad on the frame. The Silicon Micro Display is also a wearable (head-mounted) device with larger transparent displays covering both eyes. The Silicon Micro Display receives its video input via an HDMI connection.

We begin by providing an overview of our analysis, including a discussion of our fundamental research question and our human computation approach, in Section 3.1. We then turn to our experimental setup in Section 3.2 and our results in Section 3.3. We reflect on the significance of our results in Section 3.4.

3.1 Experimental Overview

Our primary goal was to assess the degree to which information on the Google Glass and Silicon Micro Display ST-1080 can be recognized or reconstructed by third parties. To make our goal more concrete, we begin with the following natural question: suppose Alice is wearing one of these displays while talking with Bob. Can Bob infer information about what Alice is doing on her device simply by looking at the light emitted from her display?

One method to attempt to answer the above question would be to conduct a study with a large number of participants playing the role of Bob. For such a study, a researcher might wear the device and do the appropriate actions to cause the device to display some specific content. Another researcher might then ask the participants to answer questions about what they see on the display.

We chose not to do the above-described study for a number of reasons, including both logistical ones (e.g., we would need to control for eye sight differences between participants, which to be accurate would require calibration on the day

of the experiments since prescriptions change over time) and practical ones (e.g., the challenge of recruiting enough participants for the study to have meaningful results). But, most importantly, we were concerned about the following: suppose that participants at a distance of (say) 1.4m from the display claim to not be able to infer any information about the contents of the display. Could we have confidence that other individuals could not do better? For example, what if a real Bob had better eye sight? Or what if a real Bob was wearing a head-mounted display of his own, with a built-in camera?

To address these concerns, we reformulated a related question: can an observer with a camera located some distance away from Alice's head-mounted display infer information about what is being shown on the display? At large distances, and with a suitable lens, the camera is likely to provide the observer with significantly more information that an observer using his or her own naked eye. Hence, the results of this experiment provide an upper bound on the information leakage from the display to an observer via the visual channel. We took multiple photos in different display and camera configurations. We then applied a human computation approach to extract information from the results: we uploaded a representative photo from each configuration to an online crowdsourcing service. We paid crowdsourced workers to answer questions about the uploaded photos. These workers act as an image processing backend.

To summarize, if the device is found to *not* leak information to an observer under our experimental setup, then we can have confidence that an observer will not be able to infer that information with his or her naked eye. This conclusion is true also when other factors are taken into consideration, such as motion caused by the person wearing the display moving his or her head.

In contrast, if our workers can infer information about the contents shown on these displays, then we have shown that there are cases where these displays are *not private* against bystanders. As we will see, this is the case – the assumption that near eye displays are private is not universally true. Future work could refine the limits of the guarantees these displays provide.

Other Scenarios. We observe that there are natural scenarios in which attackers might have capabilities similar to those embodied in our experimental setup. Namely, cameras are common in many environments (e.g., behind an ATM machine, behind a clerk at a gas station, behind one-way mirrors). Such cameras could also be used to take photos of users wearing the Glass or Silicon Micro Display, just as we did in our experiments.

3.2 Experimental Setup

We now turn to a more detailed description of our experimental setup, beginning first with the setup for obtaining photos of the Glass and Silicon Micro Displays (Section 3.2.1) and then describing our process for crowdsourcing the assessment of information leaked in the photos (Section 3.2.2).

Figure 3: Experimental setup for evaluating the display leakage of the Silicon Micro Display ST-1080.

We have put all the raw images that we collected during this study on `http://www.displayleak.com`, for reviewers to evaluate independently; for this submission, we have removed or cropped potentially de-anonymizing images.

3.2.1 Obtaining Photos

Figure 3 overviews our experimental setup.

Placement of Display. Rather than have the target display floating in free space or on a shelf, to approximate the use of the Glass or Silicon Micro Display by a person, we placed the device on a mannequin head for all of our experiments. Since the display is transparent in both directions, this setup provided a realistic background behind the display.

Distance Between Display and Camera. Psychologists have studied how far apart people are from each other when they interact. According to [21], there is the *intimate distance* region (from touching up to approximately 18 inches or 0.45 meters apart), the *personal distance* region (from approximately 0.45 meters to four feet or 1.2 meters apart), and the *social distance* region (from approximately 1.2 meters to eight feet or 2.4 meters apart). Quoting from [21], the intimate distance region is for "lovers, children and close family members;" the personal distance region is for "conversations with friends, to chat with associates, and in group discussions;" and the social distance region is for "strangers, newly formed groups, and new acquaintances."

We picked one distance in each of these regions: 0.35m, 0.7m, and 1.4m. For each experiment, we ensured that the camera lens and the outer lens of the wearable display were separated by approximately these distances. Our choice in these distances was not arbitrary. We began by picking the 0.35m value, to correspond with a distance well within the intimate distance but still at a comfortable conversation distance for people in this category. We then doubled that amount for the personal distance (0.7m), and doubled that amount again for the social distance (1.4m).

We did not take pictures of the Silicon Micro Display at 0.7m because the lens we used at 0.7m for the Glass became

unavailable to us. However, as our results suggest, bracketing the Silicon Micro Display experiments at 0.35m and 1.4m seemed sufficient for our purposes.

Camera and Camera Settings. We used a Canon Rebel XS camera (a DSLR; crop factor 1.6) to take pictures of the display. We used either a tripod or a shelf to ensure that the camera was stable and at the same height as the wearable display. We used three Canon lenses: a 24-70mm lens for the 0.35m and 0.7m Glass photos (set to 70 mm focal length), a 18-55mm lens for the 0.35m Silicon Micro Display photos (55mm focal length), and a 70-200mm lens for all 1.4m photos (200mm focal length). We set a timer on the camera so that it would delay 10 seconds before taking a picture (to prevent shaking). We also programmed the camera to use Auto Exposure Bracketing (AEB), taking three successive shots at different exposures and later choosing the best one.

Lighting. We conducted our experiments in normal building lighting conditions. Our experiments with the Glass were conducted in a room with a lighting level of between approximately 240 to 275 lux when measured at the front of the Glass. Our experiments with the Silicon Micro Display were conducted in a room with less lighting, namely a lighting level of approximately 70 lux when measured at the front of the display. The Silicon Micro Display is highly reflective, and hence that the lower lighting may have been beneficial for the experiments. Furthermore, because of the reflectivity of the Silicon Micro Display, we placed a large black poster board behind the camera when taking the photos.

Image on Display. We conducted our experiments with the devices displaying a total of four classes of images: an optometrist eye chart (both white text on a black background and black text on a white background), a password display, a Facebook-related page, and a WikiLeaks-related page. Figure 4 shows these images. We expand on our description of these images below.

We used identical images on both the Glass and the Silicon Micro Display. Namely, for the password, Facebook, and Wikileaks, we made the Glass display the appropriate image (described more below) and used the MyGlass Android application's screencast capability to display that image on a paired Android phone. We then took a screen capture of the image on the phone. Those screenshots are shown in Figure 4. We then displayed those screenshots on the Silicon Micro Display. For the eye charts, we displayed the raw eye charts on both the Glass and the Silicon Micro Display.

We use the eye charts and password displays to evaluate the text-reconstruction capability for individual letters; if we had chosen words, it might be possible for someone to reconstruct a word even if he or she cannot actually recognize every letter. The password display is from a proposed Google Glass password manager, recently described in a Communications of the ACM article on augmented reality and security [14]; the code for the password manager was available at the authors' public GitHub repository (`https://github.com/froeschele/GlassPass`). Their password manager works as follows: when the user navigates to a webpage, a modified browser will display a page-specific QR code. The user can scan the QR code with his or her Glass, and the password manager would then display to the user his or her password. Since a password is private, the privacy of this display is important to evaluate; the authors of [14] note this as a potential concern, but do not evaluate it. We obtained the eye chart from a public website [13].

Our Facebook and Wikileaks displays were obtained as follows. Using the Glass, we issued the voice command "OK, Glass" followed by "Google, Facebook" and "Google, Wikileaks." The images shown in Figure 4 are from the image that the Glass displayed in response to these voice commands. Each image includes both a logo as well as some explanatory text obtained via Wikipedia. Since application support for the Glass is still rather limited, we considered these images as the best candidates for evaluating whether a third party might be able to infer information about what activity a user might be performing with her or her wearable display. We chose Facebook as an example of a site that many people recognize, and Wikileaks as an example of an activity that some users might consider sensitive. (We observe, however, that the use of Facebook or other services in some situations might still be considered sensitive.)

3.2.2 Crowdsourced Information Leakage Assessment

We took multiple photos in each test condition: for both the Glass and the Silicon Micro Display, at distances of 0.35m, 0.7m, and 1.4m, and with each of the five images shown in Figure 4. We later manually cropped each photo to the size of the device's lens. For example, for the Glass, we cropped photos to the borders of the transparent part of the Glass (which is larger than the Glass's actual display region). Since the displays were sometimes slightly angled on the mannequin's head, the size of the resulting cropped images, while approximately the same, are not exactly the same.

We then selected the best images from each condition. In the Appendix, Figures 12, 13, and 14 respectively show the Glass images from 0.35m, 0.7m, and 1.4m away. Figures 15 and 16 respectively show the Silicon Micro Display images from 0.35m and 1.4m away. All images are shown cropped to the size of display.

While we could have tried qualitatively to assess how much information is leaked via each photo ourselves, such an assessment have been biased by our knowledge of the ground truth source images. Instead, we applied a human computation approach for extracting information from these images. We uploaded these images to Crowdflower, an online crowdsourcing service, and asked those crowdsource workers to answer questions about the uploaded images. We paid workers

Figure 4: Images displayed on the Glass and Silicon Micro Displays.

$0.50 per task.

For the eye chart and password manager, we asked participants to transcribe each line of the images. For Facebook and Wikileaks, we asked participants if they could identify the company or organization associated with each image, and then we asked participants to transcribe the text in the images.

We took a number of precautions to protect the validity of our results. For example, since there is a way for a Crowdflower worker to see the file names of the images that they are assigned, we intentionally chose to obscure the image file names so that the workers could not infer critical information about an image's contents from its file name. As another example, without appropriate precautions, a single worker could be assigned multiple eye chart transcription tasks (e.g., a single worker might be assigned both the Silicon Micro Display image at 0.35m away and the Glass image at 1.4m away). Such a situation would damage the validity of the results because the worker might remember his or her answers from one image, which would then help him or her transcribe the subsequent images. Hence, we configured the system such that a single worker would never be assigned more than one eye chart image, more than one password manager image, and more than one website image. The latter constraint was applied across both Facebook and Wikileaks, e.g., if a person was assigned a Facebook image, he or she would not be assigned a Wikileaks image. While workers could potentially see that we had three types of images available, no worker could ever see more than one image of each type.

The only exception to this constraint is that, a few weeks prior to our main crowdsourcing run, we did a trial study involving three workers and one eye chart image. Thus there is a possibility of contaminating three of our eye chart workers, but we consider this to be unlikely.

3.3 Results

Eye Chart. Table 1 presents the results from our Crowdflower-powered analysis, using seven different workers for each image. We present the raw results since we find that any textual summary alone fails to clearly capture some of the nuances of the results.

From the table, it is clear that both eye charts (white on black and black on white) can be fully reconstructed on both displays (the Glass and the Silicon Micro Display) at a distance of 0.35m.

The Silicon Micro Display is fully reconstructable at a distance of 1.4m. However, as the distances increase, the reconstructability of the Glass display decreases. Several observations arise here. First, the larger lines remained reconstructable even at greater distances. Second, the results suggest that for the Glass it is easier for the workers to reconstruct white text displayed on black background rather than black text displayed on a white background; this observation derives heavily from the workers' ability to reconstruct the white on black eye chart at 1.4m but their apparent challenges in reconstructing the black on white eye chart at the same distance. This situation seems logical since white on black is typically easier to discern for both eyes and cameras as one approaches the limits of the optical system; essentially, any problems with aberrations or over exposure will lower the contrast much faster on a mostly white image as light leaks into the thin black areas. Additionally, though perhaps unsurprisingly, we see confusion between similarly-looking characters, e.g., in several instances workers reported the "Q" character as an "O" or "G."

Password Manager. We now turn to our study of the information leakage about the contents of the password manager's display. Recall that the password manager display (shown in Figure 4) is that of a proposed Google Glass password manager from [14], with code available on GitHub. We wanted to determine whether a third party observer could reconstruct information about the displayed password, thereby violating an important security property of the password manager.

Table 2 presents the results of our Crowdflower analysis of the password manager photos. The workers were effective at reconstructing (most of) the password at a distance of 0.35m, for both the Glass and the Silicon Micro Display. For the Glass, the workers did have several errors, but as Table 2 suggests, an observer given a photo at this distance would be able to significantly reduce the search space for the password. Moreover, rather than simply ask the Crowdflower workers to transcribe the text shown in the photos, we could have asked the workers to enter each character on a separate line, along with their confidence score for each character or a sorted list of what they think each character might be. Given this information, along with general knowledge about the visual similarity between different characters in a particular font (e.g., "8" and "B" might look similar, and "1" and "l" might look similar), an adversary should be able to find the user's password much faster than via a brute-force exhaustive search.

At 1.4m, with one exception, the workers were also effective at reconstructing (most of) the password shown on the

Device	Eye Chart	Distance	Row 1 (A)	Row 2 (ZY)	Row 3 (EUWQ)	Row 4 (MNDHR)	Row 5 (EYLUZM)
Glass	White on black	0.35m	✓(7)	✓(7)	✓(7)	✓(7)	✓(7)
Glass	Black on white	0.35m	✓(7)	✓(7)	✓(7)	✓(7)	✓(7)
Glass	White on black	0.7m	✓(7)	✓(7)	✓(2), ---O(5)	✓(1), ----W, ---NK, ---NH, ----K, ----*, ---MN	S--MFH, I--NF-, I-XN-W, -DAV-P, X-ZMIH, I--BTX, ****-*
Glass	Black on white	0.7m	✓(7)	✓(7)	✓(3), ---O(2), ---G(2)	---N-(5), ----E, ---NE	---OYG, --EQ-Q, LTSG-R, L--O-Q, LT-S-R, RKHIF*, ---JIQ
Glass	White on black	1.4m	✓(7)	✓(7)	✓(6), ---O	✓(7)	✓(4), Z--M-W, ---M--, -T----
Glass	Black on white	1.4m	✓(7)	✓(7)	✓(4), ---O(2) ---G	--G-H, --OW-, -RON-, --GM-, --QN-, --OME, --ON-	******, -LEE-Z, CA-AC-, CANNON, *XADN-, HHIMH-, CABPO-
SMD	White on black	0.35m	✓(7)	✓(7)	✓(7)	✓(7)	✓(7)
SMD	Black on white	0.35m	✓(7)	✓(7)	✓(7)	✓(7)	✓(7)
SMD	White on black	1.4m	✓(7)	✓(7)	✓(7)	✓(7)	✓(7)
SMD	Black on white	1.4m	✓(7)	✓(7)	✓(7)	✓(7)	✓(7)

Table 1: Eye chart results. A "✓" means that the Crowdflower worker correctly identified all characters. A entry such as "---G" means that the worker correctly matched the characters in positions 1, 2, and 3 but had an incorrect value (in this case a "G") in the fourth position. Notation such as "---G(2)" means that two (of seven) workers had the same response. A "*" means that the character was omitted from the response (we aligned the remaining characters for optimal matching).

Device	Distance	Recognized?	Password (WpJ8swbq)
Glass	0.35m	Yes(7)	********, --+bo---, --.IBo---, ----o--g, ----Q--+, ----@--+, +-+-o--a
Glass	0.7m	Yes(0)	********, -e******, ********, ********, ********, ********, ********
Glass	1.4m	Yes(2)	********, ********, -*******, ********, ********, ********, --16w(x)
SMD	0.35m	Yes(7)	--+-----, -----*--, --------, --------, --------, --------, --------
SMD	1.4m	Yes(6)	********, --------, ---B----, ---B----, --------, ---B----, --------

Table 2: Password manager results. A "Yes(*N*)" in column 3 means that *N* of the seven workers recognized that the password manager was referring to twitter.com. For the fourth column, an entry such as --+bo--- means that the worker correctly identified the first two and last three characters, correctly identified the third character but had the wrong case (in this case, a "j" instead of a "J"), and incorrectly had "b" and "o" as the fourth and fifth characters. When a worker gave a password of an incorrect length, we picked an alignment that maximized matching. A "*" means that a character was omitted, in which case we aligned the remaining characters in an optimal manner. A full entry of "********" means that a worker did not submit an answer or submitted an answer that was clearly incorrect (e.g., an answer of "none" or "UNCLEAR").

Silicon Micro Display. The one exception was one worker who entered "none" in response to the transcription request; we hypothesis that this worker did not look closely at the image and hence only saw the reflection of the camera in the image and not the displayed text. Indeed, looking at Figure 16, the images of the Silicon Micro Display from 1.4m away is dominated by the reflection of the room.

For the Glass at 0.7m, one worker correctly observed that the photo was of a password manager display and also correctly observed the first character of the password, attempted to provide the second character (but was incorrect), and omitted the rest of the password. For the Glass at 1.4m, one worker also correctly observed that the photo was of a password manager display and correctly identified the first two characters of the password. Another worker simply entered a "W" as his or her response, and we do not know if that user correctly identified the first character and stopped or if the worker just happened to enter a "W".

Ultimately, we argue that our results suggest that information about the password is indeed leaked via the device's display even when the entire password is not reconstructable. Namely, while doing a full study of the entropy loss for the passwords is beyond the scope of this paper, and indeed not necessary to study our primary question of whether information leaks or not, visual inspection of the password manager photos and the Crowdflower responses we have received suggest that entropy does occur even at sizeable distances.

Facebook and WikiLeaks. We now turn our study of the Facebook and WikiLeaks images; the original images are shown in Figure 4. Table 3 gives the results of our Crowdflower study with the Facebook image; Table 4 gives the results of our study with the WikiLeaks image.

We first consider the Glass when displaying the Facebook image. We find that at all distances (0.35m, 0.7m, and 1.4m), workers — when asked what company or organization the image corresponded to — correctly reported Facebook. Since the text was not readable at the larger distances, however, our

Device	Distance	Recognized?	Transcription (Facebook is an online social networking service headquartered in Menlo Park, California. Its name comes from a colloquialism for the directory given to students at some America... Wikipedia)
Glass	0.35m	Yes(7)	*all-correct*(0) ; *none*(2) ; facebook is ; Facebook ; Facebook is ; a Wikipedia article ; Facebook is an online social networking service headquartered in Menlo Park, California. Its name comes from a colloquialism for the directory given to students at some American universities.[6] Facebook was founded on February 4, 2004, by Mark Zuckerberg
Glass	0.7m	Yes(7)	*all-correct*(0); *none*(5) ; f ; Something about facebook?
Glass	1.4m	Yes(6)	*all-correct*(0) ; *none*(5) ; F ; Facebook
SMD	0.35m	Yes(7)	*all-correct*(4) ; *none*(1) ; Facebook is an onlike social networking service headquartered in Menlo Park ; a wiki description of facebook
SMD	1.4m	Yes(6)	*all-correct*(3) ; *none*(0) ; F ; facebook is an ; facebook is ; facebook

Table 3: Facebook display results. A "Yes(*N*)" in column 3 means that *N* of the seven workers recognized that the image corresponded to Facebook. An "*all-correct*(*N*)" in column 4 means that *N* of the seven workers correctly transcribed all of the text at the right side of the image; we include in this count transcriptions with typos, transcriptions with the final word "Wikipedia" removed, as well as transcriptions with small deviations (e.g., the addition of the word "universities" at the end of the sentence or the addition of the word "North" before "America"). A "*none*(*N*)" in column 4 means that *N* of the seven workers failed to transcribe anything. Column 4 also lists all the transcription responses that did not fall into the *all-correct* or *none* categories.

Device	Distance	Recognized?	Transcription (WikiLeaks is an international, online, non-profit, journalistic organisation which publishes secret information, news leaks, and clssified media from anon... Wikipedia)
Glass	0.35m	Yes(0)	*all-correct*(0) ; *none*(7)
Glass	0.7m	Yes(0)	*all-correct*(0) ; *none*(7)
Glass	1.4m	Yes(0)	*all-correct*(0) ; *none*(7)
SMD	0.35m	Yes(7)	*all-correct*(6) ; *none*(0) ; WikiLeaks is an international, online, non-profit, journalistic organisation which publishes (...) information, news, leaks and classified media from (...) Wikipedia
SMD	1.4m	Yes(1)	*all-correct*(0) ; *none*(6) ; wiki leaks

Table 4: WikiLeaks display results. See the caption for Table 3 for how to interpret this table.

results suggest that the workers were cueing off the distinctive Facebook logo. This result suggest that high-level information can in some cases leak to third-party. Anecdotally, we also find that when one of us wears the Glass and accesses this same page, another of us can easily recognize the image (and its distinctive colors) from a comfortable (social) communications distance. Our observation about the workers cueing off of the Facebook logo is corroborated by the fact that none of the workers correctly identified that the WikiLeaks image corresponds to WikiLeaks, likely because the workers were not as familiar with WikiLeaks and its logo.

Another observation that we found informative is the following: one worker apparently identified that the Glass image for Facebook at 0.35m included text from a Wikipedia article and, instead of transcribing the text directly, he or she apparently went to Wikipedia and copied parts of the article directly from Wikipedia. We infer this because the worker included a fragment of a sentence that appears on the Facebook Wikipedia [20] page but does *not* appear in the image displayed on the Glass: "Facebook was founded on February 4, 2004, by Mark Zuckerberg." This incident confirms a hypothesis of ours: that when some information about the display is leaked to a third-party, the third party can use additional resources, when available, to help reconstruct the rest of the information.

At 0.35m away, workers were significantly better at reconstructing the WikiLeaks content when shown on the Silicon Micro Display than on the Glass. At 1.4m away, workers had greater success identifying and transcribing the Facebook image when shown on the Silicon Micro Display than the WikiLeaks image when shown on the same display. This is consistent with our earlier observation that workers were cueing off of the familiarity off the Facebook logo.

3.4 Discussion

We discussed some lessons from our human-powered analysis inline above. We provide further reflection here.

What should we do? Is the information leakage via the visual channel a sufficient concern to warrant defenses? Based on our results, the answer is yes, but with some clarifications. Using a good camera and lens combination, and without advanced photo-editing techniques, we find that one can infer significant information from both the Glass and the Silicon Micro Display even at social distances. The camera and

lens combination provides and upper bound for what a human might be able to reconstruct with her or her bare eye, but suggests that information does leak. The amount of information leaked depends on the content being displayed. For example, the Facebook logo is easily recognizable even at 1.4m, whereas individual text characters can be harder to reconstruct. Rather than try to formulate a model for what constitutes a dangerous amount of display leakage, we adopt the position that leaking any information can be risky since, *a priori*, we do not know all the types of content a display might render. We also recall from Section 3.1 that there are scenarios in which attackers can place cameras in the environment, e.g., at ATMs or behind one-way mirrors. Hence, our conclusion is that wearable transparent displays would greatly benefit from defenses against display leakage.

Other Side Channels. In considering these issues with wearable displays, it is important to stress that information about a user's activities may also leak in other ways. For example, returning to the scenario of Alice using the Glass and Bob as a bystander, if Alice says "OK, Glass, Google ... Facebook," then Bob could infer information about Alice's activities from both the audio channel (what Alice says, or what the Glass's speakers say back to Alice) and the visual channel (the image leaked through the transparent display). Future devices may have different user interaction modes and may leak less (or more) information via the audio channel. We believe that it is important to understand information leakage via both channels, hence our focus on the visual channel here.

A more sophisticated adversary may use other, more technical mechanisms, to infer Alice's activities. For example, a more sophisticated adversary could attempt to intercept wireless traffic coming from Alice's wearable display or the paired phone or computer. An even more sophisticated adversary may try to analyze the electromagnetic emanations from the device, as pioneered in the public literature with works such as [9, 10, 19]. Our goal was not to analyze the capabilities of such an adversary, and indeed we argue that the likelihood of such adversaries manifesting in practice is probably less than the likelihood of Bob simply trying to look at Alice's Glass and see what he can see, or of another adversary hiding a camera behind a one-way mirror. Nevertheless, we encourage future works in analyzing the electromagnetic emanations from wearable displays like the Glass.

Comparing the Glass and the Silicon Micro Display. We found numerous differences between the Glass and the Silicon Micro Display. There are the obvious structural differences — the Glass display being smaller than the Silicon Micro Display. During our human-powered analysis, we found the Silicon Micro Display to be significantly more reflective than the Glass. The Glass seems to let more light through the display than the Silicon Micro Display. In our setup, focusing on the eye chart, we found that white text on a black background was easier to see on the Glass than black text on a white background. The opposite seems to be the case of the Silicon Micro Display. Namely as we discussed above in the context of the password manager application, the white on black password manager image for the Silicon Micro Display was hard for one worker to see at 1.4m (but all the other workers were able to reconstruct most of the password, likely because they looked closer at the image and saw the faint text). See Figure 17 for images of a black on white password manager screen as shown on the Silicon Micro Display; this photo has more visible content than the white on black password manager images in Figures 15 and 16.

Figure 5: Display leakage from the Lumus. The image is clear and visible from a distance, but only a small portion is visible from each vantage point.

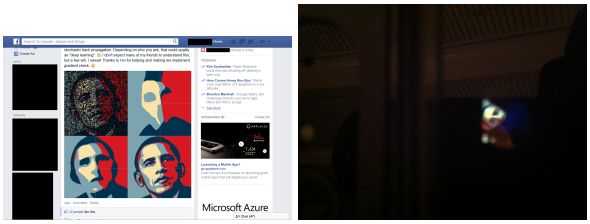

Figure 6: On the left, a picture of a web page we showed on the Meta One. We have redacted potentially identifying pieces of the web page. On the right, a photograph of the Meta One from the "outside," with part of the web page clearly visible.

Ultimately, our assessment is that the Silicon Micro Display leaks more information than the Glass when reflection is not an issue (e.g., when the third party is viewing the display directly rather than at an angle or when the room is dark). We make one additional observation, however. Namely, while we use the same images on the Glass and Silicon Micro Display in order to directly compare them, the Silicon Micro Display could display much higher resolution content to the user, and the higher resolution content may be harder for a third party to visually reconstruct. On the other hand, the LCoS display panel used in the SMD display is significantly larger than that in the Glass, so the angular resolution is comparable.

The Lumus and the Epson Moverio. We also obtained access to two other transparent wearable displays: the Epson Moverio and the Lumus. We did not study these in detail but

make the following observations. We found that the Lumus device also leaks information to parties on the other side of the display; see Figure 5. However, we found that only a small part of the image is visible from each vantage point. We found that the Moverio did not leak visual information to a third party. We discuss this further in the next section.

The Meta One. Finally, we examined the Meta One development kit, which consists of a head mounted display with integrated depth camera. The Meta One is an early prototype and developer kit, so we stress that the final product may not have the same optical properties. We found that the Meta One leaks information in a similar fashion as the Lumus. Figure 6 shows on the top a screen shot of a web page we showed on the display, and on the bottom a photo taken from the outward facing side of the Meta One. While only a small part of the image is visible from each vantage point, moving the camera allowed us to clearly see different parts of the image.

4 Root Causes

We now turn to a discussion of how head-mounted displays work and, in particular, how current designs impact display leakage. While we believe that this discussion is new to the computer security community, it is based in large part on one of our experiences in optics research, available public information on these devices [8], and our inferences from examining the devices. Our perspective is different than that of traditional display designers: namely, optical engineers typically optimize display properties like resolution, brightness, and contrast [8] whereas our focus here is on the properties of the devices that lead to display leakage. We stress that different display designs make different tradeoffs, and because we believe display leakage has not been a priority for designers, therefore the presence or absence of display leakage in shipping displays is incidental to their other characteristics.

Head-mounted displays designed for augmented reality (AR HMDs) often leak light to the outside world due to the transparent and symmetrical nature of the display system. Typically they include some combination of OLED or LCoS microdisplay, a projector and/or imaging optical element such as a lens or mirror to produce a virtual image at some point distant to the viewer (typically between 1.6m and infinity), and either free space or a transparent medium for both the image light and the light from the outside world to travel through. The designs of AR HMDs can be simple or complex; we will only discuss a few simple variants which illustrate our points.

Birdbath displays. One common design for AR HMD optical systems are variations of the classic birdbath system, where a spherical reflector is used to form the image in conjunction with a beamsplitter that also allows light to enter from the outside world. Google Glass is a well-known implementation of this, embedded in a solid block which provides a robust, simple optical system. The reflective imaging element is at the end of the block and the beamsplitter is embedded within it as shown in Figure 7.

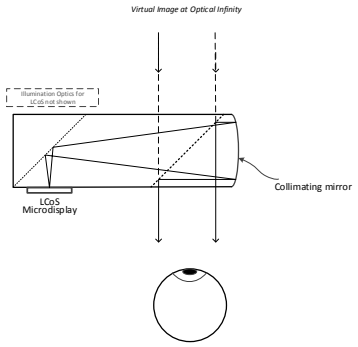

Figure 7: The optical pathway for Google Glass, which is an example of a "birdbath" design.

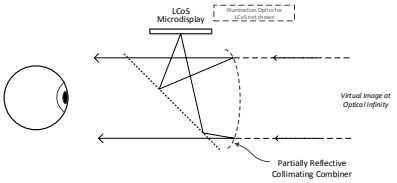

Figure 8: The optical pathway for the Silicon Micro Display ST-1080, another "birdbath" design.

Unfortunately, this type of system with a simple beamsplitter reflects a fair amount of light directly emitted from the microdisplay directly out of the display to the outside world. Our experiments with the Glass in Section 3 confirm this behavior. This may have been a conscious decision by the system designers, to alert people in the area that Glass is on and potentially recording them. However, as we have seen above, even the unmagnified image can effective spill confidential information to a properly equipped adversary.

The Silicon Micro Display is another example of a birdbath system. The Silicon Micro Display places the reflector between the viewer and outside world, which affords a larger field-of-view (FOV). However, for maximum brightness the SMD uses a very reflective combiner and only allows roughly 10% of the outside light in, also limiting the amount of light leakage. It might thus be considered more as a VR display with a bit of ambient awareness for safety or comfort reasons.

Waveguide systems. Another type of AR optical system is a waveguide where light from the microdisplay is collimated (image formed at infinity), effectively forming a tiny projection system. The collimation light is directed into a thin transparent substrate where it can bounce around a number of times by total internal reflection (TIR) before being ejected by some features intentionally placed inside or on top of the

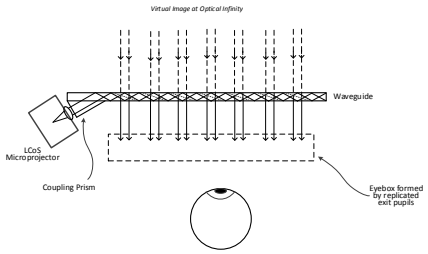

Figure 9: The optical pathway for a Lumus DK-40, an example of a "waveguide" design.

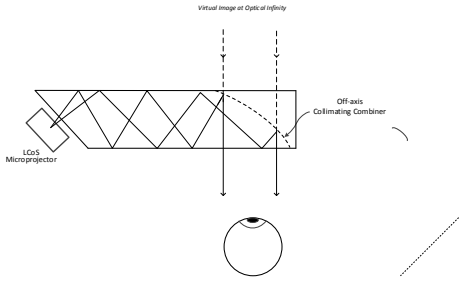

Figure 10: The optical pathway for the Epson Moverio BT-200.

waveguide. Typically there are a number of places where the light is being ejected, each of which forms a bundle of rays through which the image can be seen. The replication of these bundles form multiple exit pupils, or places where the pupil of the eye can see the image. By forming multiple exit pupils as a larger eyebox, less or no adjustment is needed for comfortable viewing by people with different sized heads and inter-ocular spacing.

Perhaps the simplest example of such a system is made by Lumus, as seen in Figure 9. Progressively more reflective beamsplitter bands are used to eject consistent amounts of projector light with successive bounces of the image path, allowing a relatively large eyebox and a wider FOV than Glass. While these beamsplitter use sophisticated polarization coatings, a fair amount of light (several percent) still manages to escape through the wrong side of the waveguide to the outside world. Even worse, this light forms replicated, spatially separated, collimated and magnified images of the microdisplay that can be seen from far away.

The only mitigating factor is that only a small part of the image can be seen from a given vantage point at a given time. Figure 5 shows an example of a Lumus devices display leakage from a single vantage point. However, we conjecture that a sophisticated adversary could rapidly collect a large sample of images using a hidden video camera as the wearer moves their head, and stitch together the images off-line.

A newer system by Optivent uses a Fresnel-like reflector embedded within to eject the light out, instead of the flat tilted beamsplitter elements of the Lumus [11]. While we have not yet seen nor evaluated this system, it is reasonable to assume that the faceted reflector will at least partially scramble any image light that leaks to the outside world.

Alternate Approaches. The Epson Moverio uses a variant not quite like either of these. In the Moverio, the image path is folded via TIR within a (thick) waveguide-like cavity but the pupil is not replicated as in a waveguide. Instead, an off-axis reflector is used to both form the image and reflect the light into the eye. Any light that passes through the reflector continues reflecting through the waveguide and is not leaked out. The width of the cavity and the size of the reflector determine the exit pupil size, and the burden of the designer is more in creating a high-quality image with an off-axis optical system. We do not know whether Epson intentionally chose this design to avoid leaking visual information to third parties or if it is a byproduct of some other goal; but we would have expected to see public discussions of this property if it was indeed a design goal.

5 Defenses

We describe two basic defensive strategies to reduce display leakage, and prototype one of them.

Polarization. Our first approach builds on the properties of polarized light. In particular, the ambient and display light can be separated by polarization so that any light from the display with one polarization will be blocked from leaving the display by a polarizer. Thus, we can prevent light leaving the display away from the user from being reconstructed into a coherent image.

This approach means that 55-60% of outside light will also be blocked, dimming the outside world in a manner comparable to untinted polarized sunglasses. We argue that this may be an acceptable tradeoff for providing a private display, especially for form factors such as Google Glass that do not cover the user's entire field of view.

Alternatively, a polarizing privacy filter may be used only selectively when the user is in non-private environments or viewing sensitive content, using a small sticker-like filter. This is similar to how privacy filters are used with laptops today. With the appropriate operating system and device support, applications could choose not to show sensitive data (or to prompt the user before showing it) when the filter is not present. For example, a hardware switch could be depressed when the filter snaps into place, signaling the filter's presence to the operating system, which could surface this information to applications. Adding an absorbing polarizer to a wearable device adds a negligible cost (a few dollars).

Narrowband. Another general means of defense would be to use narrowband (i.e., laser) illumination of an LCoS microdisplay and use narrowband filters to prevent the display light from being seen. If the light and filters are narrowband

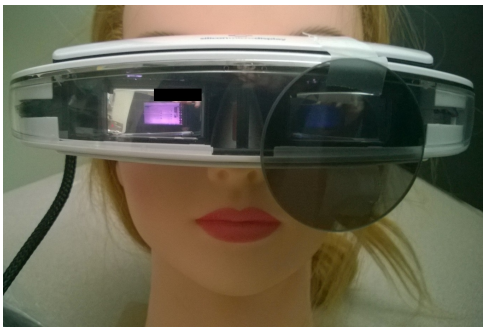

Figure 11: The polarization defense applied to the Silicon Micro Display ST-1080. We have modified the figure because an author was visible in the reflection of one eyepiece.

enough, this would allow greater than 50% transmission of outside light through the system. Any ambient light in the spectral bandwidth of the display reflected from the surface of the filters back into the environment will also decrease the S/N ratio of the leaked signal. Finally, a film that efficiently polarizes over 3 narrow spectral bands for Red, Green and Blue could be combined with narrowband illumination of an LCoS display. This would provide the highest level of transmission possible while blocking the display light leakage.

Design. Changing the illumination to a narrowband system requires extensive engineering costs and should ideally be undertaken at the early stages of the design. We decided instead to use a common absorbing polarizer to illustrate what can be done — perhaps even as an aftermarket solution — to address the problem of display leakage in an economical manner. Furthermore, the majority of AR HMDs use Liquid Crystal on Silicon (LCoS) as their microdisplay elements, and the resulting image light is inherently polarized.

Implementation. We intended to use Google Glass but ran into an unexpected complication. While publicly available teardowns of Glass show that a LCoS device is used, the solid glass structure is highly birefringent. This leads to unpredictable variation in the polarization of the image light as it is leaked out away from the viewer by the beamsplitter, rendering the defense ineffective. It may be difficult to fabricate the solid beamsplitter structure without stress-induced birefringence.

The Silicon Micro Devices system has little if any birefringence. The display is readily adaptable to the proposed polarization solution. Figure 11 shows a comparison between display leakage for a modified and an unmodified eyepiece on our Silicon Micro Devices display.

Evaluation. Visual inspection and our photographs show the difference is quite evident. Anyone could simply tape polarizing film to the outside of their SMD device experience vastly lower display leakage, albeit with 60% reduction of outside light. Note that the SMD uses a highly reflective combiner element; simply using a less reflective coating could compensate for this with a minor penalty of display brightness. To illustrate, let us assume the coating transmits 8% of the outside light and reflects 90% of the display light. Changing the coating to transmit 18% of the outside world light would mean reflecting 80% of the display light, or only about a 12% reduction from the current design.

6 Conclusions

Multiple researchers have assumed transparent near-eye displays are private. tO capture this, we introduced the goal of *display privacy*. This goal captures adversaries should not learn display contents from light that they leak to the outside.

We proposed ways to measure *display leakage*. We found that two shipping displays, the Google Glass and the Silicon Micro Display ST-1080, both suffer from a concerning amount of display leakage. We also observed display leakage in the Meta One and Lumus displays.

We proposed an inexpensive defense based on polarization. Surprisingly, this defense does not work for the Google Glass due to the glass being highly birefringent; this defense does work with the unmodified Silicon Micro Display. We also showed how alternative optical pathways yield better display privacy through less light leaked to the outside world. We lay the foundation for future displays to combat display leakage and obtain better display privacy.

References

[1] R. T. Azuma. A survey of augmented reality. *Presence: Teleoperators and Virtual Environments*, 6(4):355–385, Aug. 1997.

[2] M. Backes, T. Chen, M. Duermuth, H. Lensch, and M. Welk. Tempest in a teapot: Compromising reflections revisited. In *IEEE Symposium on Security and Privacy*, 2009.

[3] M. Backes, M. Duermuth, and D. Unruh. Compromising reflections -or- how to read LCD monitors around the corner. In *IEEE Symp. on Security and Privacy*, 2006.

[4] M. Corporation. Computer filters, 2014. `http://solutions.3m.com/wps/portal/3M/en_US/ergonomics/home/products/computerfilters/`.

[5] A. Greenberg. Google Glass snoopers can steal your passcode with a glance, 2014. `http://www.wired.com/2014/06/google-glass-snoopers-can-steal-your-passcode-with-a-glance/`.

[6] S. Jana, D. Molnar, A. Moshchuk, A. Dunn, B. Livshits, H. J. Wang, and E. Ofek. Enabling fine-grained permissions for augmented reality applications with recognizers. In *USENIX Security Symposium*, 2013.

[7] S. Jana, A. Narayanan, and V. Shmatikov. A Scanner Darkly: Protecting User Privacy from Perceptual Applications. In *IEEE Symposium on Security and Privacy*, 2013.

[8] B. Kress and T. Starner. A review of head-mounted displays (HMD) technologies and applications for consumer electronics. In *Photonic Applications for Aerospace, Commercial, and Harsh Environments IV*, volume 8720 of *Proceedings of SPIE*, 2013.

[9] M. G. Kuhn. Electromagnetic eavesdropping risks of flat-panel displays. In *Workshop on Privacy Enhancing Technologies*, 2004.

[10] M. G. Kuhn and R. J. Anderson. Soft tempest: Hidden data transmission using electromagnetic emanations. In *Information Hiding*, 1998.

[11] K. Mirza and K. Sarayeddine. Key challenges to affordable see through wearable displays: The missing link for mobile AR mass deployment, 2014. `http://optinvent.com/HUD-HMD-benchmark`.

[12] E. Ofek, S. T. Iqbal, and K. Strauss. Reducing disruption from subtle information delivery during a conversation: mode and bandwidth investigation. In *CHI*, 2013.

[13] optometrial.com. Snellen vision chart, 2014. `http://www.optometrial.com/snellen-vision-chart-downloadable-graphic-free`.

[14] F. Roesner, T. Kohno, and D. Molnar. Security and Privacy for Augmented Reality Systems. *Communications of the ACM*, 57:88–96, 2014.

[15] F. Roesner, D. Molnar, A. Moshchuk, T. Kohno, and H. J. Wang. World-driven access control for continuous sensing. In *ACM Conference on Computer and Communications Security*, November 2014.

[16] M. Simkin, A. Bulling, M. Fritz, and D. Schroeder. Ubic: Bridging the gap between digital cryptography and the physical world. In *ESORICS*, 2014.

[17] I. E. Sutherland. A head-mounted three dimensional display. In *Proceedings of the Fall Joint Computer Conference, Part I*. ACM, 1968.

[18] R. Templeman, M. Korayem, D. Crandall, and A. Kapadia. PlaceAvoider: Steering first-person cameras away from sensitive spaces. In *Network and Distributed System Security Symposium (NDSS)*, 2014.

[19] W. van Eck. Electromagnetic radiation from video display units: An eavesdropping risk? *Computers & Security*, 1985.

[20] Wikipedia. Facebook, 2014. `http://en.wikipedia.org/wiki/Facebook`.

[21] Wikipedia. Personal Space, 2014. `http://en.wikipedia.org/wiki/Personal_space`.

[22] Wikipedia. Shoulder surfing, 2014. `http://en.wikipedia.org/wiki/Shoulder_surfing_(computer_security)`.

[23] C. Winkler, J. Gugenheimer, A. de Luca, G. Haas, P. Speidel, D. Dobbelstein, and E. Rukzio. Glass unlock: Enhancing security of smartphone unlocking through leveraging a private near-eye display. In *CHI 2015 Notes.*, 2015.

[24] D. Yadav, B. Ionascu, S. Ongole, A. Roy, and N. Memon. Design and analysis of shoulder surfing resistant PIN based authentication mechanisms on google glass. In *Financial Cryptography*, 2015.

A Raw Pictures Provided To Crowdflower

Figure 12: Images of the Glass from 0.35m away. The reader may wish to experiment by zooming into these images in the PDF.

Figure 13: Images of the Glass from 0.7m away. The reader may wish to experiment by zooming into these images in the PDF.

Figure 14: Images of the Glass from 1.4m away. The reader may wish to experiment by zooming into these images in the PDF.

Figure 15: Images of the Silicon Micro Display from 0.35m away. The reader may wish to experiment by zooming into these images in the PDF.

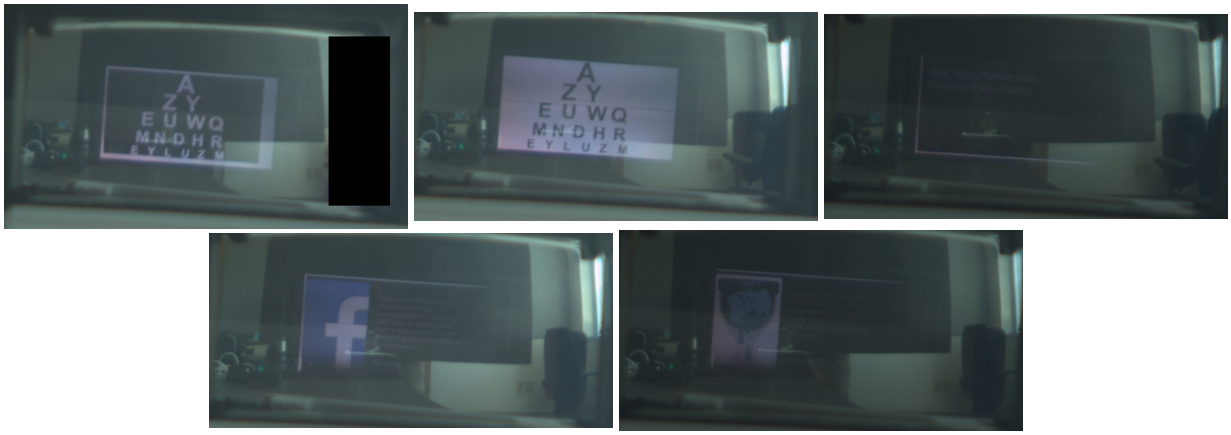

Figure 16: Images of the Silicon Micro Display from 1.4m away. The reader may wish to experiment by zooming into these images in the PDF. We have placed a black box over part of the first image to anonymize the submission; part of one of the authors is visible in the original image.

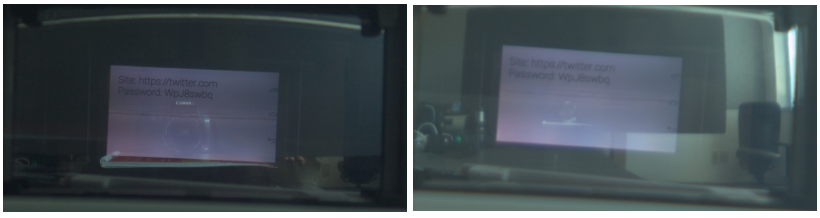

Figure 17: Images of the Silicon Micro Display, displaying a black on white password manager screen, from 0.35m and 1.4m away. The reader may wish to experiment by zooming into these images in the PDF.