

Security and Privacy for Augmented Reality Systems

Franziska Roesner
University of Washington

Tadayoshi Kohno
University of Washington

David Molnar
Microsoft Research

1 Introduction

Augmented reality (AR) technologies promise to enhance our perception of and interaction with the real world. Unlike virtual reality systems, which replace the real world with a simulated one, augmented reality systems sense properties of the physical world and overlay computer-generated visual, audio, and haptic signals onto real-world feedback in real time. In this article, we consider the security and privacy concerns associated with AR systems themselves as well as those that arise from the supporting technologies.

Researchers have explored the idea of augmented reality since the 1960s, when Sutherland described a transparent head-mounted display showing three-dimensional information [33]. Since the 1990s, AR as a research area has focused on overcoming challenges with display technology, tracking and registration to properly align virtual and real objects, user interfaces and human factors, auxiliary sensing devices, and the design of novel AR applications [1, 2, 6, 22, 36, 41].

However, it is only recently that early-generation AR technologies have begun shipping commercially. For example, Google recently released a limited number of its Google Glass, heads-up glasses for augmented reality applications. Many other early-generation AR applications are enabled by the ubiquity of smartphones and other mobile devices. Examples include the Word Lens iPhone application — an application that overlays translated text on the camera’s view of foreign text — and Layar, a geolocation-based AR platform that allows developers to create augmented reality layers for the world (e.g., for game playing); see Figure 1. The recent advent of 1 GHz processors, location sensors, and high resolution, autofocus cameras in mobile phones has made these applications possible.

In this article, we take a broad view of the AR space, considering both direct applications of AR as well the technologies necessary to support these applications. Beyond the mobile phone, devices are becoming available that enhance sensing, display, and data sharing, which will enable more complex AR systems. For example, Looxcie — an over-the-ear, always-on video camera — includes a feature enabling wearers to share their live video feed with anyone else in the world. Microsoft’s



Figure 1: **Phone-Based Augmented Reality.** On the left, a picture of Word Lens, an iPhone application that provides seamless “in-picture” translation (source: <http://www.flickr.com/photos/neven/5269418871/>). Here the app translates the word “Craft” from English to Spanish and then back again. On the right, a picture of Layar, an “augmented reality browser” shipping on Android phones (source: <http://site.layar.com/company/blog/make-your-own-layar-screen-shot-with-the-dreamcatcher/>).




Figure 2: **Wearable Input and Output.** On the left, a Looxcie body-worn camera worn by a ranger in Kenya (source: <http://looxcie.com/index.php/image-gallery>). On the right, a Google Glass prototype in June 2012 (source: <http://www.flickr.com/photos/azugaldia/7457645618>).

SDK for Kinect [20], which provides accurate motion sensing by combining an RGB camera, a depth camera, and a multi-array microphone, has enabled numerous prototype AR applications. In addition to Google Glass, transparent, wearable displays are now available for research purposes from several companies, such as Vuzix, Lumus, and Meta SpaceGlasses. Figure 2 shows examples of such input and output devices. See Appendix A for a summary of AR-enabling technologies; many of these technologies are shipping today, while others are still experimental.

These technologies will enable commercial augmented reality applications and are at the cusp of significant innovation, which will bring significant benefits to many users. However, these technologies may also bring

unforeseen computer security and privacy risks. Previous research in the AR space has rarely considered these issues. Rather than wait for these technologies to fully mature and then retroactively try to develop security and privacy safeguards, we argue that now is the time to consider security and privacy issues, while the technologies are still young and malleable. To guide this process, we ask the following questions: (1) What new security and privacy research challenges arise with AR systems and the technologies that support them? (2) What novel opportunities do AR technologies create for improving security and privacy?

We find that AR technologies form an important, new, and fertile playground for computer security and privacy research and industry. Of course, these technologies should leverage standard security best practices, such as on-device and network encryption. Nevertheless, we find unique obstacles — such as handling conflicts between multiple applications sharing an AR system’s output — that are simultaneously intellectually challenging yet surmountable. Other challenges, such as access control for data, are well known in other arenas but become even more important for AR technologies with their always-on, always-sensing inputs. Given the future importance of AR technologies, researchers already tackling these issues in other domains can find value in refocusing their attention on AR applications.

In addition to presenting new challenges, AR systems present opportunities for new applications that improve security and privacy. For example, these technologies can provide personal digital views of content on personal displays. Imagine a password manager that superimposes visual indicators over the correct keys for a complex password when a user looks at a keyboard, or an application that alerts the user when someone is lying.

We explore new security and privacy challenges presented by AR technologies in Section 2, defensive directions in Section 3, and new applications of AR systems to known security and privacy issues in Section 4.

2 Challenges

The AR applications and technologies that we consider in this article may have any or all of the following characteristics, in addition to the traditional definition of aligning real and virtual objects in real-time:

- A complex set of input devices and sensors that are always on (e.g., camera, GPS, microphone).
- Multiple output devices (e.g., display, earpiece).
- A platform that can run multiple applications simultaneously.
- The ability to communicate wirelessly with other AR systems.

In this section, we present a set of security and privacy challenges that come with these novel technologies and

their applications, as summarized in Figure 3. We organize these challenges along two axes: system scope and functionality. On one axis, we consider AR systems of increasing scope: single applications, multiple applications within a single AR platform, and multiple communicating AR systems. The challenges in each category first appear at that level of system complexity. For each scope, we further categorize challenges as related to input, output, or data access. We encourage future designers of AR technologies to consider security and privacy challenges along both axes.

Readers familiar with smartphone security may observe some overlap between those challenges and the set that we present here. We note that some techniques from smartphone security may be applicable to AR technologies; others will need to be rethought in this new context. We return to this discussion in Section 3.

2.1 Challenges with Single Applications

We first consider threats and challenges limited in scope to a single AR application.

Output. Users must place significant trust in AR applications that overlay real-world visual, auditory, or haptic perceptions with virtual feedback. Devices providing immersive feedback can be used by malicious applications to *deceive users about the real world*. For example, a future malicious application might overlay an incorrect speed limit on top of a real speed limit sign (or place a fake sign where there is none), or intentionally provide an incorrect translation for real-world text in a foreign language. More generally, such an application can trick users into falsely believing that certain objects are or are not present in the real world.

Malicious applications can use similar techniques to cause *sensory overload* for users. By flashing bright lights in the display, playing loud sounds, or delivering intense haptic feedback, applications could physically harm users. Such attacks are not unprecedented: attackers have targeted epilepsy forums, posting messages containing flashing animated gifs to trigger headaches or seizures [24]. Emerging AR platforms must consider and prevent these types of attacks.

These output attacks are more serious in immersive AR applications than they are in today’s desktop or handheld computing scenarios both because it is harder for users to distinguish virtual from real feedback and because it may be more difficult for users to remove or shut down the system. As a last resort for output attacks, users must be able to easily and reliably return to the real world, i.e., with all output devices verifiably turned off.

In the near term, removing the system is a simple way to achieve this *return to reality*. However, future wearable systems may be hard or impossible for users to remove (e.g., contact lenses [23] or implanted devices),

	Single Application	Multiple Applications	Multiple Systems
<i>Output</i>	Deception attacks Overload attacks Trusted path to reality	Handling conflicts Clickjacking	Conflicting views
<i>Input</i>	Input validation	Resolving focus	Aggregate input
<i>Data Access</i>	Access control for sensor data Bystander privacy	Cross-app sharing	Cross-system sharing

Figure 3: **Security and Privacy Challenges for AR Technologies.** We categorize these challenges by two axes: challenges related to output, input, and data access, as arise in single applications, multi-application systems, and multiple interacting systems.

and today’s non-wearable systems may already be hard for users to evade. For example, several automotive manufacturers have produced windshields that display augmented content over the user’s view of the road [5]. In these cases, the system should have a trusted path for the user to return to reality, analogous to Ctrl-Alt-Del on Windows computers. Determining the best such sequence, or the right input mode (e.g., gestures or speech), requires research for each AR system. Another approach may be to reserve a trusted region of the display that always shows the real world.

Input. Augmented reality applications will undoubtedly face similar *input validation and sanitization* challenges as conventional applications. For example, a translation application that parses text in the real world may be exploited by maliciously crafted text on a sign. Traditional input validation techniques are likely to apply, but the designers of AR systems should be aware of their necessity in this new context.

Data Access. To provide their intended functionality, AR applications may require *access to a variety of sensor data*, including video and audio feeds, GPS data, temperature, accelerometer readings, and more. As in desktop and smartphone operating systems, an important challenge for AR systems will be to balance the access required for functionality with the risk of an application stealing data or misusing that access. For example, a malicious application may leak the user’s location or video feed to its backend servers. The existing proof-of-concept PlaceRaider attack [34] shows that smartphone sensors can be used to gather enough information to create three-dimensional models of indoor environments.

Unlike most of today’s desktop and smartphone applications, complex AR applications will require rich, always-on sensing. For example, an application that automatically detects and scans QR codes requires constant access to video stream data, as does an application that automatically detects when the user is entering a password on another device and provides password assistance (see Section 4). As a result, these privacy risks are much greater than in conventional systems.

AR systems should take approaches that limit these

risks. For example, individual applications will likely not need access to all sensor data. Perhaps an application only requires access to a portion of the screen when the user is in a certain location, or only needs to know about certain objects that the system recognizes (e.g., via the Kinect’s skeleton recognizer), rather than needing access to the entire raw camera feed. AR system designers must consider the appropriate granularity for these permissions, and the design of usable permission management interfaces will be important. Existing manifest or prompt-based solutions as used in smartphones are unlikely to scale in a usable way, and the long-term (rather than one-time) data access needs of AR applications make the application of in-context access control solutions like user-driven access control [28] not straightforward.

Always-on cameras and other sensors will also create a *privacy risk for bystanders*, which Krevelen and Poelman identify as a challenge for widespread social acceptance of AR [36]. Bystanders should be able to opt out of or be anonymized (e.g., blurred) in the recordings of others; prior work has examined such issues [9, 31]. AR users may need methods to prove to skeptical bystanders that such safeguards are in place. Legislation or market forces may lead to cameras that respond to requests from other devices or the environment; news reports suggest that Apple has considered adding such a capability to the iPhone to prevent videotaping of live events, such as concerts [4]. Cameras may also alert bystanders while recording, such as by flashing a light [36] or by providing access to more complex policy information [19].

The CVDazzle project [10] pursues a different approach — using makeup to confuse face detection algorithms — that provides privacy without compliant cameras. The key limitation is that CVDazzle is painstakingly hand-tuned for one particular face detection algorithm. A research question is to find a general algorithm for synthesizing makeup that fools face detection.

2.2 Challenges with Multiple Applications

Though AR applications are often conceived and prototyped in isolation, we can expect that future AR platforms, like those built on Google Glass or the Microsoft



Figure 4: **Multi-Application Augmented Reality.** Emerging and future augmented reality platforms will support multiple applications running simultaneously, sharing input and output devices, and exposing data and APIs to each other. In a multi-application AR system, applications like those depicted in this mockup will share output devices, including displays, audio output, and haptic feedback. Conflicts among these applications can result in security concerns (Section 2.2).

Kinect, will support multiple applications running simultaneously, sharing input and output devices, and exposing data and APIs to each other (Figure 4). Researchers must anticipate these developments and ensure that an “operating system for augmented reality” is designed with appropriate considerations for security and privacy.

Output. In a multi-application AR system, applications will share output devices, including displays, audio output, and haptic feedback. Conflicts among multiple applications attempting to use these output devices can lead to security concerns. For example, a malicious application might try to obscure content presented by another application (e.g., visually or aurally covering up a correct translation with an incorrect one).

Nevertheless, output sharing will be necessary to provide desirable functionality in AR systems. For example, a user may wish to simultaneously view content overlaid on their view of reality from multiple applications, such as directions supplied by a maps application, a social feed summarizing the activity of nearby friends, the track currently playing in a music application, and so on. Thus, the naive solution, in which only one application controls the display at a time (as in Android today, for instance), is insufficient.

Thus, future AR systems must *handle conflicts* between multiple applications attempting to produce output. For example, five applications may all want to annotate the same object (e.g., with a translation subtitle), and the system will need to prioritize them. It may furthermore be important for users to know which content was generated by which application—for instance, whether an annotated product recommendation comes

from a friend or an advertiser. AR system designers must create interfaces that make the origins of displayed content clear to or easily discoverable by users.

Traditional attacks based on the manipulation of output may require new approaches or new formulations in the AR context. For example, in today’s systems, applications can mount *clickjacking attacks* that trick users into clicking on sensitive user interface elements from another application (e.g., to post something on the user’s social media profile). These attacks generally work either by manipulating the display of the sensitive element—by making it transparent or partially obscuring it in a clever way—or by suddenly displaying sensitive elements just before users click in a predictable place. Future applications on AR systems may develop new techniques for tricking users into interacting with elements, and system designers must anticipate these threats. For example, an AR application could attempt to trick a user into interacting with an object in the physical, rather than the virtual, world.

Input. Users will likely not interact with AR systems using traditional input methods like clicking on a mouse or even using a touchscreen. Instead, users may increasingly interact with these systems using subtle input to haptic sensors (e.g., embedded in gloves), using voice, or with the aid of gaze tracking technologies. With these input techniques and multiple running applications, it will be nontrivial for the system to *resolve which application is in focus* and should thus receive input.

For example, today’s voice interactions happen either following an explicit user action indicating the destination application (e.g., clicking on the “Siri” button on an

iPhone) or on systems in which only one application can ever receive voice input (e.g., on the Xbox). When multiple applications are active and might receive voice or other input at any given time, there must be either a usable way for users to bring applications into focus, or for the system to determine the correct intended destination for input commands when focus is ambiguous. We emphasize that future AR systems are likely to run multiple applications simultaneously, many of them running and listening for input without having any visible output. Improperly designed focus resolution may make it easy for malicious applications to steal user input intended for another application (e.g., to steal a password intended for the login box of another application). For example, a malicious application may attempt to register a similar-sounding verbal keyword as another, sensitive application, intentionally increasing input ambiguity.

Data Access. As in traditional operating systems, AR applications will likely wish to expose APIs to each other, and users may wish to share virtual objects between applications. Researchers must explore appropriate *access control models for cross-application sharing*. Certainly lessons from traditional access control design can be applied in this space, but new technologies and environments may require new approaches. For example, copy-and-paste and drag-and-drop are established user gestures for sharing data between traditional applications and thus have access control implications. A long line of work in desktop and smartphone systems has attempted to map user actions to application privileges (examples include [21] and [28]); AR systems will need to evolve new user gestures to indicate sharing intent. Additionally, AR systems are unlikely to display applications in labeled windows the way that traditional desktop operating systems do, so new interaction paradigms will be needed to enable users to identify applications and indicate which application should receive shared data.

2.3 Challenges with Multiple Systems

Moving beyond a single AR system running multiple applications, we consider the interactions between multiple AR systems belonging to different users. Prior work in AR proposes collaborative applications among multiple users of an AR system. These applications include multi-player games [11, 32, 40], telepresence for remote conferencing [16], and face-to-face collaboration [26]. These types of applications pose additional security and privacy challenges.

Output. Multiple users may have *differing views of the world* presented by their respective AR systems. For example, different users may see different virtual advertisements superimposed on real-world billboards, or different users watching a presentation may be shown different content based on their access levels (i.e., one user may

see top-secret footnotes while others do not). Such conflicting views will require users to manage mental models of who can perceive which information, lest they accidentally reveal private information intended only for themselves. Addressing this concern will require innovations in interface design for aiding users in this task.

Input. A rise in the complexity of AR systems and applications will be tightly coupled with a rise in the number and complexity of sensor inputs provided by enabling technologies. This abundance of sensor input from many users will in turn lead to novel collaborative sensing applications, which can themselves feed data back into AR applications. For example, Google already uses data collected by users' smartphones to estimate traffic conditions, which is then reported back to user's phones [8]. This type of data is necessary to enable future AR applications displayed on the car's windshield, for example.

However, this type of *aggregate input* can be used by malicious users to fool the data collection systems. For example, a review site might leverage location tracking to measure a restaurant's popularity by noting the average number of people present during the day. A canny restaurateur may then pay people to stand in the restaurant without buying anything. The restaurant's measured popularity rises but has no relationship to its quality.

AR technologies that constantly collect data will drive the adoption of such collaborative sensing applications; thus, these security concerns will increase in importance. As another example, the Community Seismic Network aggregates accelerometer sensor data of many individuals to detect and predict earthquakes; an attacker could manipulate the sensors to "spoof" unusual seismic activity, e.g., by encouraging many individuals monitored by the project to jump at once in the context of an unrelated game. (For example, Improv Everything [13] asks users to play provided audio files as a designated time and follow the audio instructions.) Trusted sensors [30]—while important to prevent other attacks—do not help in these cases, as real-world conditions are manipulated.

Data Access. In addition to showing different content to different users, communicating AR systems will *allow users to share virtual content* with each other. For example, one user may create a virtual document within their private AR system and later choose to share its display with the systems of other users. Some sharing may even be implicit; imagine an AR system that automatically uses the camera feeds of nearby users to provide a given user with a real-time 3D model of him or herself.

The implicit or explicit sharing of data across separate AR systems can enable valuable applications. However, appropriate access control models and interfaces are needed to allow users to manage this sharing. Today, users already have difficulty forming mental models

of their privacy settings on services like Facebook because of the complexity of relationships between people and data items [18]. The vast amount of data collected by AR systems and the integration of virtual objects with the real world will make this problem only more difficult.

3 Defensive Directions

In this section, we outline several defensive directions for AR technologies. First, some of the security and privacy challenges associated with AR technologies are similar to those faced by smartphones today, such as the privacy of sensor data and cross-application sharing. In some cases, an appropriate defensive direction for augmented reality is to *adapt smartphone solutions*. For example, permission manifests and the app store review process may be adopted in the short term.

In the long term, however, there are several reasons that approaches in the AR context must differ from smartphone solutions. First, an analysis of the resource needs of smartphone applications [28] showed that most require only one-time or short-term access to most resources, making solutions that require in-context user interactions (such as user-driven access control [28]) feasible. By contrast, AR applications will require long-term or permanent access to sensor data at a scale beyond smartphone applications. Further, AR resource access will not be as clear to users and to bystanders as in the smartphone context—for example, an AR system’s camera will always be on, whereas a smartphone’s camera, even if turned on by malware, provides much less data while the phone is in the user’s pocket. Thus, we argue that it is important to consider full-fledged future AR contexts when designing solutions in this space.

Along these lines, *new research into AR-specific solutions* will be needed. For example, researchers have begun considering operating system support specific to AR [7]. AR applications—and the underlying OS—naturally follow the pipeline shown in Figure 5, so research can be characterized accordingly, and different research models can assume different boundaries between the application and the OS. In the first stage, sensing, an application (or the OS) gathers raw sensory data such as audio, video, or radio waves; research here includes limiting which sensed information is collected (e.g., polite cameras [9, 31]) or limiting on its use (e.g., retention policies). Second, in the recognition stage, machine learning algorithms extract objects with high-level semantics: as an example, the figure shows a Kinect skeleton, a face, the associated name, and voice command triggers. Related research includes changing objects to cause false negatives (e.g. CVDazzle [10]) and policies governing application access to objects [15]. Finally, the application (or the OS) renders on top of the user’s senses, such as vision and hearing. Research here




Figure 5: **AR Pipeline.** AR applications (1) gather sensory data, from which they (2) extract objects with high-level semantics. Finally, they (3) render on top of the user’s senses. (Source of Google Glass image on the right: <http://www.flickr.com/photos/azugaldia/7457645618>).

includes uncovering invariants that must be respected to avoid harming the user and building a performant “trusted renderer” that respects these invariants.

Not all defensive directions for augmented reality will consist of technical solutions. Instead, some challenges may call for *social, policy, or legal approaches*. For example, in Section 2.1 we discussed potential policies for bystander opt-outs and compliant cameras. Other issues will similarly benefit from non-technical approaches.

Finally, we call for an *augmented reality testbed* for researchers working in this space. Most experimental AR applications today rely on the Microsoft Kinect or smartphone platforms like Layar; both involve only single applications running at one time, thereby hiding challenges that arise as AR systems increase in complexity.

4 Novel Applications

Though augmented reality technologies create important security and privacy concerns, there is also an unexplored opportunity for them to enhance security and privacy through their application to existing problems. In this section, we consider opportunities for new security and privacy enhancing applications enabled by AR technologies and systems. Our list is undoubtedly incomplete; we hope to see rich future work in this area.

4.1 Leveraging Personal Views

AR systems that integrate heads-up or other personal displays (like Google Glass) can leverage these personal views to address existing security and privacy concerns—in particular, protecting private data and improving password management.

Personal displays present a strong *defense against shoulder surfing*, as users may interact with applications visible in their own view only. For example, someone using a laptop on an airplane today exposes everything they view and type to their seat neighbors, and researchers have demonstrated that footage from low-cost cameras can be used to reconstruct a user’s typing on a virtual mobile keyboard [25]. A personal heads-up display combined with a haptic sensor for discreet input would allow



Figure 6: **Prototype AR Password Manager.** Our Chrome extension (background) displays a QR code representing the current website. In response to the voice command “find password,” our Google Glass application (foreground) scans this QR code and displays the stored password for that site privately in the heads-up display.

for greatly improved privacy.¹

Personal displays further enable *encrypted content in the real world* that can only be decrypted by the AR systems of the intended recipients. For example, a company can post encrypted notices on a bulletin board that employees can read through their company-issued AR systems, but that visitors to the company’s building cannot. (Storing only the key, not the encrypted content, on a server accessible by the AR system requires adversaries to find the physical notice rather than simply compromise the company’s servers.) Precursors of such a system are possible today using smartphones and 2D barcodes that encode URLs to data with appropriate access control; augmented heads-up displays will prevent the need for a manual scan.

AR systems can also act as an enhanced *password manager* for users, presenting passwords or password hints via the personal display. For example, a display could outline the appropriate characters that the user must enter on legacy devices like ATM PIN pads. Users could then be assigned strong passwords, as they would never need to actually remember them. This application requires markerless tracking and a system design that properly protects the stored passwords.

As a concrete example, we have implemented a prototype password manager application consisting of a Google Glass application and a browser (Chrome) extension (see Figure 6). The Chrome extension modifies the browser’s UI to display a QR code representing the

¹We observe that see-through displays, such as that used by Google Glass, may not be fully private from external observers. For example, images taken of the display using a telephoto lens may be used to reconstruct the content of the screen, similar to reconstructing content from screen reflections [3]. Future research should fully characterize this threat and design appropriate defenses.

website currently displayed to the user (the website in the browser’s address bar). Users can ask the Google Glass application to scan these QR codes and consult its password database by using the voice command “OK Glass, find password.” If the user has previously stored a password for that website, the application displays the password; otherwise, the user can enroll a new password by asking the Chrome extension to generate an enrollment QR code and asking the Glass to store the new password using the “enroll password” voice command. We have made the code for our prototype available at <https://github.com/froeschele/GlassPass>.

By designing the QR code displayed by the browser extension to include a secret shared between the browser and the phone, this application could furthermore serve as *phishing protection*, as websites would not be able to create and display forged QR codes that would map to legitimate passwords in the password manager.

4.2 Leveraging Complex Sensor Systems

AR systems benefit from the combination of multiple input and sensing devices, which can be combined to enhance digital and physical security, privacy, and safety.

Future systems can leverage AR technologies to *detect privacy or security conditions* of which the user should be alerted. For example, rather than relying on compliant cameras to shield users from unwanted recording, a system could alert users when it detects camera lenses pointed at them, using (for instance) computer vision to detect the glint of light reflecting off a lens [35]. It could also detect some forms of eavesdropping, e.g., a laser microphone pointed at a window.

Such systems could also *detect physical deception attempts*. For example, an AR system could estimate the size and shape of an ATM card slot, then issue a warning if it appears a card skimming device has been added. Similarly, existing work on computerized interpretation of facial expressions [12] could be applied to behavior-based lie detection [38]. One of our colleagues refers to this application as “spidey sense.”

Beyond storing passwords, AR systems can be used for *implicit authentication* of their users. The plethora of sensors attached to people using these technologies can be used to authenticate them with biometric and behavioral characteristics. Prior work has examined the possibility of such mechanisms on mobile phones [14, 27]; AR systems would provide far more powerful authentication. Similarly, sensor data could be used to help with authorization and access control decisions.

Beyond the sensors attached to an individual (e.g., Alice), the sensors of bystanders could also be used to authenticate her by providing the authentication system with third-party visual, audio, and other sensory views of Alice. This *third-party authentication* system would

distribute trust to systems and persons with no incentive to falsely authenticate Alice.

5 Conclusion

Augmented reality systems, with their sophisticated and pervasive input, output, and processing capabilities, have the potential to significantly benefit many users. To complement ongoing innovations in AR technologies, we argue that now is also the time to define a roadmap for protecting the computer security and privacy of AR systems—before these systems become widely deployed and their architectures become entrenched. To catalyze this roadmap, we consider new security and privacy challenges posed by these systems, and we explore opportunities afforded by these technologies to create novel privacy- and security-enhancing applications.

Acknowledgements

This work is supported in part by the National Science Foundation (Grants CNS-0846065, CNS-0905384, and a Graduate Research Fellowship under Grant DGE-0718124) and by a Microsoft Research Fellowship. We thank Luis Ceze, Lydia Chilton, Alexei Czeskis, Tamara Denning, Karl Koscher, Brandon Lucia, Alex Moshchuk, and Helen Wang for early conversations, as well as Nicki Dell, Brandon Lucia, Bryan Parno, Karin Strauss, and anonymous reviewers for feedback on earlier versions.

References

- [1] AZUMA, R., BAILLOT, Y., BEHRINGER, R., FEINER, S., JULIER, S., AND MACINTYRE, B. Recent advances in augmented reality. *IEEE Computer Graphics and Applications* 21, 6 (2001), 34–47.
- [2] AZUMA, R. T. A survey of augmented reality. *Presence: Teleoperators and Virtual Environments* 6 (1997), 355–385.
- [3] BACKES, M., CHEN, T., DUERMUTH, M., LENSCH, H., AND WELK, M. Tempest in a Teapot: Compromising Reflections Revisited. In *IEEE Symposium on Security and Privacy* (2009).
- [4] BUSINESS INSIDER. This apple patent will shut down your camera at live concerts, 2011. <http://www.businessinsider.com/iphone-concert-patent-2011-6>.
- [5] CNN. Augmented-reality windshields and the future of driving, 2012. <http://virtual.vtt.fi/virtual/proj2/multimedia/alvar.html>.
- [6] COSTANZA, E., KUNZ, A., AND FJELD, M. Human machine interaction. Springer-Verlag, 2009, ch. Mixed Reality: A Survey, pp. 47–68.
- [7] D’ANTONI, L., DUNN, A., JANA, S., KOHNO, T., LIVSHITS, B., MOLNAR, D., MOSHCHUK, A., OFEK, E., ROESNER, F., SAPONAS, S., VEANES, M., AND WANG, H. J. Operating system support for augmented reality applications. In *USENIX Workshop on Hot Topics in Operating Systems* (2013).
- [8] GOOGLE. Crowdsourcing road congestion data, 2009. <http://googleblog.blogspot.com/2009/08/bright-side-of-sitting-in-traffic.html>.
- [9] HALDERMAN, J. A., WATERS, B., AND FELTEN, E. W. Privacy Management for Portable Recording Devices. In *Proceedings of the 3rd ACM Workshop on Privacy in Electronic Society* (2004).
- [10] HARVEY, A. CVDazzle: Camouflage from Computer Vision. <http://cvdazzle.com/>.
- [11] HENRYSSON, A., BILLINGHURST, M., AND OLLILA, M. Face to face collaborative ar on mobile phones. In *4th IEEE/ACM International Symposium on Mixed & Augmented Reality* (2005).
- [12] HOQUE, M. E., MCDUFF, D., AND PICARD, R. W. Exploring temporal patterns in classifying frustrated and delighted smiles. *IEEE Transactions on Affective Computing*, 3 (2012), 323–334.
- [13] IMPROV EVERYWHERE. The Mp3 Experiments, 2012. <http://improveverywhere.com/missions/the-mp3-experiments/>.
- [14] JAKOBSSON, M., SHI, E., GOLLE, P., AND CHOW, R. Implicit Authentication for Mobile Devices. In *4th USENIX Workshop on Hot Topics in Security (HotSec)* (2009), USENIX.
- [15] JANA, S., MOLNAR, D., MOSHCHUK, A., DUNN, A., LIVSHITS, B., WANG, H. J., AND OFEK, E. Enabling fine-grained permissions for augmented reality applications with recognizers. Tech. Rep. MSR-TR-2013-11, Microsoft Research, February 2013.
- [16] KATO, H., AND BILLINGHURST, M. Marker tracking and hmd calibration for a video-based augmented reality conferencing system. In *IEEE/ACM Workshop on Augmented Reality* (1999).
- [17] LAYCOCK, S., AND DAY, A. A survey of haptic rendering techniques. *Computer Graphics Forum* 26, 1 (2007), 50–65.
- [18] MADEJSKI, M., JOHNSON, M., AND BELLOVIN, S. M. The Failure of Online Social Network Privacy Settings. Tech. Rep. CUCS-010-11, Dept. of Comp. Science, Columbia University, 2011.
- [19] MAGANIS, G., JUNG, J., KOHNO, T., SHETH, A., AND WETHERALL, D. Sensor Tricorder: What does that sensor know about me? In *Proceedings of the 12th Workshop on Mobile Computing Systems and Applications (HotMobile)* (2011), ACM.
- [20] MICROSOFT. Kinect for Windows, 2012. <http://www.microsoft.com/en-us/kinectforwindows/>.
- [21] MILLER, M. S. *Robust Composition: Towards a Unified Approach to Access Control and Concurrency Control*. PhD thesis, Johns Hopkins University, Baltimore, MD, USA, 2006.
- [22] PAPAGIANNAKIS, G., SINGH, G., AND MAGNENAT-THALMANN, N. A survey of mobile and wireless technologies for augmented reality systems. *Computer Animation and Virtual Worlds* 19 (2008), 3–22.
- [23] PARVIZ, B. For your eye only. *IEEE Spectrum* 46 (2009), 36–41.
- [24] POULSEN, K. Hackers assault epilepsy patients via computer. WIRED Magazine, 2008. <http://www.wired.com/politics/security/news/2008/03/epilepsy>.
- [25] RAGURAM, R., WHITE, A. M., GOSWAMI, D., MONROSE, F., AND FRAHM, J.-M. iSpy: automatic reconstruction of typed input from compromising reflections. In *18th ACM conference on Computer and Communications Security*.
- [26] REITMAYR, G., AND SCHMALSTIEG, D. Mobile collaborative augmented reality. In *4th International Symposium on Augmented Reality* (2001).
- [27] RIVA, O., QIN, C., STRAUSS, K., AND LYMBERPOULOS, D. Progressive authentication: deciding when to authenticate on mobile phones. In *21st USENIX Security Symposium* (2012).
- [28] ROESNER, F., KOHNO, T., MOSHCHUK, A., PARNO, B., WANG, H. J., AND COWAN, C. User-Driven Access Control: Rethinking Permission Granting in Modern Operating Systems. In *IEEE Symposium on Security and Privacy* (2012).
- [29] SAPONAS, T. S., TAN, D. S., MORRIS, D., BALAKRISHNAN, R., TURNER, J., AND LANDAY, J. A. Enabling always-available input with muscle-computer interfaces. In *22nd ACM Symposium on User Interface Software and Technology* (2009).
- [30] SAROIU, S., AND WOLMAN, A. I am a sensor, and I approve this message. In *Proceedings of the 11th Workshop on Mobile*

- Computing Systems and Applications (HotMobile)* (2010), ACM.
- [31] SCHIFF, J., MEINGAST, M., MULLIGAN, D. K., SASTRY, S., AND GOLDBERG, K. Y. Respectful Cameras: Detecting Visual Markers in Real-Time to Address Privacy Concerns. In *Int'l Conference on Intelligent Robots and Systems (IROS)* (2007).
 - [32] STARNER, T., LEIBE, B., SINGLETARY, B., AND PAIR, J. Mind-warping: Towards creating a compelling collaborative augmented reality game. In *ACM Intelligent User Interfaces* (2000).
 - [33] SUTHERLAND, I. E. A head-mounted three-dimensional display. In *Fall Joint Computer Conference, American Federation of Information Processing Societies* (1968).
 - [34] TEMPLEMAN, R., RAHMAN, Z., CRANDALL, D. J., AND KAPADIA, A. Placeraider: Virtual theft in physical spaces with smartphones. *CoRR abs/1209.5982* (2012).
 - [35] TRUONG, K., PATEL, S., SUMMET, J., AND ABOWD, G. Preventing camera recording by designing a capture-resistant environment. In *Ubicomp* (2005).
 - [36] VAN KREVELEN, D., AND POELMAN, R. A survey of augmented reality technologies, applications, and limitations. *The International Journal of Virtual Reality* 9 (2010), 1–20.
 - [37] VIOLA, P., AND JONES, M. Robust Real-time Object Detection. In *International Journal of Computer Vision* (Hingham, MA, USA, 2004), vol. 57, pp. 137–154.
 - [38] VRIJ, A., EDWARD, K., ROBERTS, K., AND BULL, R. Detecting deceit via analysis of verbal and nonverbal behavior. *Journal of Nonverbal Behavior* 24 (2000), 239–263.
 - [39] VTT TECHNICAL RESEARCH CENTRE OF FINLAND. ALVAR Software Library, 2009. <http://cnn.com/2012/01/13/tech/innovation/ces-future-driving/>.
 - [40] WAGNER, D., PINTARIC, T., LEDERMANN, F., AND SCHMALSTIEG, D. Towards massively multi-user augmented reality on handheld devices. In *3rd International Conference on Pervasive Computing* (2005).
 - [41] ZHOU, F., DUH, H. B.-L., AND BILLINGHURST, M. Trends in augmented reality tracking, interaction and display: a review of ten years of ISMAR. In *7th IEEE/ACM International Symposium on Mixed and Augmented Reality* (2008).

A Appendix: Technologies

The table below summarizes commercial and experimental AR-enabling technologies. The table categorizes technologies based on their availability. Some technologies are commercially available “off the shelf” today. Others are under development and exist on an experimental basis.

	Commercially Available Today	Experimentally Only
<i>Sensors (Inputs)</i>	Body-worn RGB cameras GPS (error of 5 meters or more) Accurate motion sensing (e.g., Kinect)	Haptic sensors [29]
<i>Feedback (Outputs)</i>	Opaque near-eye display Phone display/speaker Invisible Bluetooth earpiece	Transparent near-eye display Embedded displays (e.g., contact lenses [23]) Haptic feedback [17]
<i>Services</i>	Simple cloud services (e.g., photo gallery) Marker-based tracking [39] Good face detection (not recognition) [37] Expensive or cheap but inaccurate transcription	Complex cloud services (e.g., object recognition) Markerless tracking Good face recognition Cheap accurate transcription
<i>Sharing</i>	Selective sharing (photos, videos, location)	Automatic sharing