Augmented Reality: Hard Problems of Law and Policy

Franziska Roesner
UW Tech Policy Lab and
Computer Science & Engineering
University of Washington
Seattle, WA USA
franzi@cs.washington.edu

Tamara Denning
UW Tech Policy Lab and
Computer Science & Engineering
University of Washington
Seattle, WA USA
tdenning@cs.washington.edu

Bryce Clayton Newell
UW Tech Policy Lab and the
Information School
University of Washington
Seattle, WA USA
bcnewell@uw.edu

Tadayoshi Kohno
UW Tech Policy Lab and
Computer Science & Engineering
University of Washington
Seattle, WA USA
yoshi@cs.washington.edu

Ryan Calo
UW Tech Policy Lab and the
School of Law
University of Washington
Seattle, WA USA
rcalo@uw.edu

Abstract

Augmented reality (AR) technologies are poised to enter the commercial mainstream. Using an interdisciplinary research team, we describe our vision of AR and explore the unique and difficult problems AR presents for law and policy—including around privacy, free speech, discrimination, and safety.

Author Keywords

Augmented reality; privacy; usability; law; policy

ACM Classification Keywords

H.5.1—Information interfaces and presentation (e.g., HCI): Artificial, augmented, and virtual realities.
K.5.0—Legal aspects of computing: General.

Introduction

Although a vision for augmented reality (AR) has been around for decades, the technology is today poised to enter the mainstream. Household name technology companies are investing billions in AR in one form or another. AR holds enormous promise to enhance a wide spectrum of human activities—economic, cultural, and social. Of course, the technology can also compromise a variety of human values, and has the potential to alter our society in ways prosaic and profound.

Figure 1: An example of AR functionality within a single application on a traditional platform, WordLens does real-time in-image translation. (Image copyright: Neven Mrgan. Image source: http://www.flickr.com/photos/neven/5269418871/)
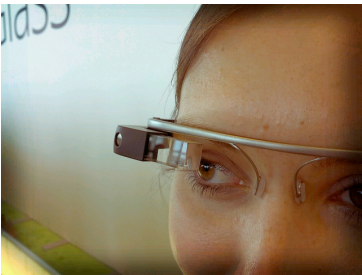

Figure 2: An example of a dedicated AR platform, Google Glass. While Glass displays mostly static content, other emerging technologies, such as Epson Moverio, promise to provide more integrated AR experiences. (Image copyright: Antonio Zugaldia. Image source: http://www.flickr.com/photos/azugaldia/7457645618)

We represent an interdisciplinary team of technologists and legal experts interested in the impact of emerging technology on society. What follows is our initial examination of AR from a technical and legal perspective. Our hope for this paper is to highlight the various difficult and unique problems AR presents and to begin to examine how the technical and legal community may engage with them.

AR raises many of the same concerns—around privacy, for instance, or security—as previous or constituent technologies such as the mobile phone. We focus here on what tends to distinguish AR and the "hard problems" we believe AR will present. Specifically, we focus on AR's capacity for always-on recording of everything a user hears and sees ("input"), as well as AR's ability to overlay information on top of physical reality ("output"). These relatively unique aspects of AR raise distinct issues that strain prevailing conceptions of privacy, free speech, discrimination, and tort law.

The paper proceeds as follows. The first part provides a technical overview of AR so as to be clear about our mental model of the technology and to delineate the scope of our research. The second part leverages this description to discuss the unique legal issues AR may raise. A final part provides our preliminary takeaways and solicits feedback from workshop participants. Did we describe AR accurately? Are the issues we raise the correct ones? And have we identified the proper strategies to address these issues?

Technical Overview

Any legal discussion must be grounded in an agreed upon understanding of the artifact under consideration. We therefore first provide a brief technical overview of augmented reality, distinguishing salient technical features and design parameters.

Definition

Augmented reality is an approximate class of technologies; definitions of AR vary, and there are not exact criteria that determine whether or not a system qualifies as AR. In general, however, AR applications and/or systems will have most or all of the following properties. They may:

- Sense properties about the real world.
- Process in real time.
- Output information to the user, including via visual, audio, and haptic [7] means, often overlaid on the user's perception of the real world.
- Provide contextual information.
- Recognize and track real-world objects.
- Be mobile or wearable.

We note that many of the above properties also individually apply to technologies that are not themselves AR technologies. For example, continuously sensing wearable devices—like the GoPro wearable camera—share several of the properties above, but do not properly constitute AR systems. In general, when a non-AR technology shares some of the AR properties listed above, it will also share a subset of the consequences and implications. For example, both AR glasses and the GoPro wearable camera raise concerns about audiovisual recording and bystander privacy.

Figure 3: An example of an AR application on Xbox, Double Fine Happy Action Theater. As evidenced by this example, AR technologies are not necessarily wearable. (Image copyright: Double Fine Productions / Microsoft Game Studios. Image source: http://marketplace.xbox.com/en-US/Product/Double-Fine-Happy-Action-Theater/66acd000-77fe-1000-9115-d802584111f3)

Examples

Augmented reality functionality can be implemented (1) within a single application running on a more traditional platform (e.g., the WordLens [10] and Layar [8] smartphone applications, or games for Xbox Kinect), or (2) within a full-fledged or dedicated augmented reality system (such as Google Glass [6], Meta SpaceGlasses [9], and Epson Moverio [4]). Figures 1-3 show examples.

These instantiations are not fundamentally different: the behavior of applications and the capabilities that they utilize are what qualify them (or not) as AR. Although many existing examples provide only limited AR functionality, we anticipate that the underlying technology will advance in the near future.

Design Parameters

Figure 4 shows an overview of an augmented reality system or application. We consider two primary technical design parameters: (1) whether AR functionality is implemented by the operating system or third-party applications, and (2) whether processing is done on the AR device itself or elsewhere on a server. These design parameters raise a number of issues with both technical and legal implications, including:

- First vs. third-party innovation: Where functionality is implemented presents a tension between control of the ecosystem and the freedom to innovate.
- Jailbreaking and reflashing: The operating system of an AR device cannot necessarily be trusted.
- Performance and battery life: Remote processing and rendering may save battery life at the expense of performance (due to network latency).
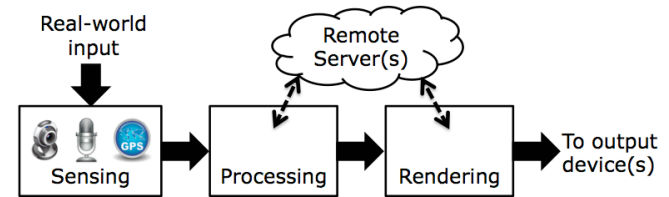

Figure 4. Overview of an augmented reality application or system. Depending on the implementation, each action (stage of the pipeline) may be performed by the operating system, by an application, or both. The operating system and applications may choose to offload some or all processing and rendering onto remote servers.

- Security and privacy on remote servers: Concerns may arise depending on who owns, administers, and can obtain legal or illegal access to remote servers and communication networks.
- Crowdsourcing for processing: In addition to improved computer vision, AR systems may supplement with expert or non-expert processing by crowdsourced workers (e.g., [3]).

Legal Issues

Law has its own logic, grouping concepts together for reasons that are sometimes historical or otherwise idiosyncratic. Here we organize our analysis around the properties of AR itself—specifically, AR's capacity for both inputting (sensing) and outputting (displaying) information, each of which raises distinct if related policy problems. In the first section, we discuss privacy and other issues primarily relating to input. In the second section, we discuss competition, bias, reliance, and other issues primarily relating to output. We note the overlap throughout.

**Case Study #1: Input
AR & the First Amendment**

In recent years, federal courts in the United States have begun to recognize a First Amendment right to record or photograph matters of public interest, including the activities of police officers and other public officials (e.g., [5]). These cases highlight an emerging right to access and document information, spurred by legal challenges to state laws that prohibit recording audio (and video, in some cases) without obtaining the consent of the subjects. AR potentially complicates this situation because a person using an AR device to continuously sense and record information may record both matters of public and purely private interest. If AR records continuously without distinguishing between subjects (e.g., a police officer or member of the public), it is unclear how far a First Amendment right to gather information could be stretched without eviscerating important privacy protections.

Input Issues

AR raises privacy concerns practically by definition. The technology relies upon the prospect of recording and analyzing the physical world in or near real time. AR thus represents a particularly acute example of a broader trend toward information collection, processing, and dissemination. This capacity—as well as the responses we might anticipate by the public and the state—also implicates values such as ownership and free speech. In brief, our research surfaced the following issues related to collection:

- Today the courts treat nearly any expectation of privacy in public as unreasonable. But technologies such as GPS and drones that are capable of widespread or constant surveillance at low cost are testing the limits of this doctrine [12]. AR will put additional pressure on this cracking edifice because it has the potential to record persistently, source and present related information from various sources to users, and blend seamlessly into the environment.

- American constitutional law also assumes no reasonable expectation of privacy in information conveyed to a third party [11]. AR has the potential to convey one's entire stream of observation to a company for analysis and storage, with unclear constitutional import. Design choices about whether to store data locally or in the cloud (or to provide user with a choice) directly affect the level of legal privacy protections afforded that information vis-à-vis the user.

- Historically, free speech interests have involved the right to express oneself in various media. AR tests the limits of a burgeoning free speech right, recognized by a handful of courts (e.g., [5]), to photograph public

officials or matters of public interest. Case study #1 (see sidebar) details this issue.

- AR complicates intellectual property law by gathering and potentially transforming copyrighted or trademarked material that appears in the real world. For example, recording copyrighted material likely constitutes copying, for purposes of copyright infringement, at the moment of capture—as well as when copies are saved to external (temporary or permanent) storage. Of course, the usual defenses to infringement (e.g., fair use) apply in these scenarios, but the potentially pervasive and persistent sensing of copyrighted material by AR technologies, combined with manipulation or output issues, raises difficult new questions about how existing intellectual property law will apply to new situations made possible my AR.

- The form factor of recording equipment has an effect on rulings in the legal landscape. In areas where there is a reasonable expectation of privacy, the presence of obvious recording equipment—like a shouldercam—is considered to serve as a cue that recording may be taking place. While early AR rigs—such as those worn by Steve Mann or Thad Starner—were fairly obvious, modern AR systems are leaning towards more inconspicuous form factors. This, in turn, can have an effect on legal rulings regarding captured footage.

- AR systems might also be designed to allow remote environmental triggers to control when sensing capabilities should be disallowed (for example, a movie theater may limit the ability of devices to record while a movie is screening) [1]. This possibility raises novel questions about limits on First Amendment information gathering rights, device ownership, intellectual property protections, and personal privacy.

Output Issues

AR not only gathers information about the world, these systems also overlay information on top of it. This capacity raises a host of independent issues that are in many cases unique to or amplified with the technology. In brief, our research surfaced the following issues related to display:

- Users of AR may rely upon data that leads to their injury or other harm. Information provided to AR users may be false, incomplete, or misleading. Scenarios range from obscuring a road sign or distracting a driver, to misidentifying a plant or mushroom as safe to eat or failing to inform a user when a potentially dangerous situation is sensed by the technology. This capacity will test the limits of product liability law, among other areas, and the specific design of these systems (e.g., whether they are designed for specific or general purposes) may alter the legal outcomes.

- AR can furnish users with truthful information they should not have, or at least that they cannot legally use to make decisions. Thus, for instance, a system could use facial recognition to pull up a job candidate's mug shot, social media profile, or relationship status in a jurisdiction that does not permit employers to discriminate based on arrest history, marital status, or other information that may be available through technological intervention. Thus, the use of AR could contribute to forms of illegal discrimination, raising possible legal liability for users and developers. Case study #2 (see sidebar) details this issue.

- AR could even prove the source of a new category of "digital assault," i.e., intentional interference with an AR user to cause fear or other harm. Tort law purports to cover such transgressions, but there are next to no test cases to date. There are, however, preliminary examples—for instance, hacking a website for epileptics to attempt to induce seizures, or advertising for exterminator services by creating the illusion that a spider has run across the user's screen [2]. These factors suggest that the use of AR to surprise, scare, or harm an AR user (particularly when the technology can sense the user is in a vulnerable situation; for example, while driving a car or when the person is depressed or unhappy) may lead to potential liability for something akin to digital assault.

Preliminary Takeaways

Finally, we highlight several key takeaways for technologists and legal scholars considering augmented reality:

- AR raises a variety of difficult legal questions by virtue of its twin abilities to sense (input) information from the real world and overlay (output) information to users. This distinction between input and output is meaningful and helps clarify the legal implications of the use of AR in society.

- Input issues deal with the potentially constant, passive collection of information and include privacy in public, intellectual property infringement, and free speech rights to record.

- Output issues, on the other hand, deal with overlaying information over physical reality and include frightening and distracting the user, providing incorrect information, or providing information the user may not lawful or ethically use to make a decision.

- Certain legal issues occur almost irrespective of AR's design or technical architecture. Other legal issues become more or less salient depending on how AR is

implemented. Designers should consider, for instance, where data is stored, what security measures are implemented, whether and how recording is indicated to bystanders, and whether the platform will be responsive to external commands to shut down sensors.

We hope to continue our discussions and begin to develop a set of concrete recommendations for addressing these issues. We solicit feedback from readers and workshop participants about these issues as well as any we may have failed to raise.

Conclusion

AR is rapidly entering the mainstream. It resembles previous consumer technologies but has key differences that may pose unique and difficult challenges for society. Particular implementations of AR strain prevailing conceptions of privacy and free speech, and have the potential to compromise the user by overlaying information on the world that is erroneous, dangerous, or legally problematic. We canvassed those challenges here in an effort to identify the specific ways technologist and policymakers can and should engage with AR going forward and to solicit commentary on our initial conclusions.

Acknowledgements

References

[1]   Apple. Apparatus and methods for enforcement of policies upon a wireless device. U.S. Patent 8254902 (filed 2008).

[2]   Calo, R. When Apps Attack. In *Forbes.com* (November 17, 2013).

[3]   Deng, J., Krause, J., and Fei-Fei, L. Fine-Grained Crowdsourcing for Fine-Grained Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition* (2013).

[4]   Epson Moverio. http://www.epson.com/cgi-bin/Store/jsp/Landing/moverio-bt-200-smart-glasses.do

[5]   Glik v. Cunniffe, 655 F.3d 78 (1st Cir. 2011).

[6]   Google Glass. https://glass.google.com/

[7]   Gupta, S., Morris, D., Patel S.N., Tan, D. Airwave: Non-Contact Haptic Feedback Using Air Vortex Rings. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing* (2013).

[8]   Layar. https://www.layar.com/

[9]   Meta SpaceGlasses. https://spaceglasses.com/

[10] QuestVisual: WordLens. http://questvisual.com/

[11] Smith v. Maryland, 442 U.S. 735 (1979).

[12] United States v. Jones, 132 S.Ct. 945 (2012).