

AUGMENTED REALITY

A Technology and Policy Primer



TECH POLICY LAB

UNIVERSITY *of* WASHINGTON

AUGMENTED REALITY: A TECHNOLOGY AND POLICY PRIMER

Tech Policy Lab, University of Washington, September 2015

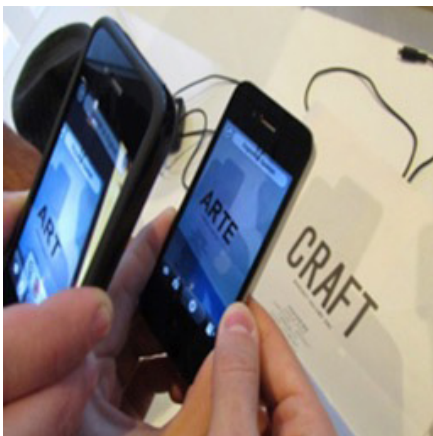
Visitors to the Haunted Mansion ride at Disneyland may recall the moment when, passing by a long mirror, ghostly figures appeared riding alongside them in the cart. This effect was an early and fun example of augmented reality (AR), a set of technologies that overlay information onto everyday experience.

The vision for AR dates back at least until the 1960s with the work of Ivan Sutherland. In a way, AR represents a natural evolution of information communication technology. Our phones, cars, and other devices are increasingly reactive to the world around us. But AR also represents a serious departure from the way people have perceived data for most of human history: a Neolithic cave painting or book operates like a laptop insofar as each presents information to the user in a way that is external to her and separate from her present reality. By contrast, AR begins to collapse millennia of distinction between display and environment.



Disneyland guests encounter a "ghost" in the Haunted Mansion. (Image source: www.disney.com)

Today, a number of companies are investing heavily in AR and beginning to deploy consumer-facing devices and applications. These systems have the potential to deliver enormous value, including to populations with limited physical or other resources. Applications include hands-free instruction and training, language translation, obstacle avoidance, advertising, gaming, museum tours, and much more.



Using the Word Lens app to translate the word "Craft" in real-time. (Image source: <http://www.flickr.com/photos/neven/5269418871/>)



Across Air app displaying information about the New York subway. (Image source: www.acrossair.com)



Navdy HUD combining its turn-by-turn directions and telephone functions. (Image source: <https://www.navdy.com/press>)

AR also presents novel or acute challenges for technologists and policymakers, including *privacy*, *distraction*, and *discrimination*.

This whitepaper—which grows out of research conducted across three units through the University of Washington’s interdisciplinary Tech Policy Lab—is aimed at identifying some of the major legal and policy issues AR may present as a novel technology, and outlines some conditional recommendations to help address those issues. Our key findings include:

1. AR exists in a variety of configurations, but in general, AR is a mobile or embedded technology that senses, processes, and outputs data in real-time, recognizes and tracks real-world objects, and provides contextual information by supplementing or replacing human senses.
2. AR systems will raise legal and policy issues in roughly two categories: collection and display. Issues tend to include privacy, free speech, and intellectual property as well as novel forms of distraction and discrimination.
3. We recommend that policymakers—broadly defined—engage in diverse stakeholder analysis, threat modeling, and risk assessment processes. We recommend that they pay particular attention to: a) the fact that adversaries succeed when systems fail to anticipate behaviors; and that, b) not all stakeholders experience AR the same way.
4. Architectural/design decisions—such as whether AR systems are open or closed, whether data is ephemeral or stored, where data is processed, and so on—will each have policy consequences that vary by stakeholder.

The whitepaper follows a method developed by our Lab for examining new technologies. The method, which also provides a roadmap for the whitepaper, consists of the following elements:

We work with technologists—in this case, computer science professors and students—to define the technology we are examining as precisely as possible.

We look to the humanities and social sciences—here, information science—to think through the impact of the technology on various stakeholders. Mindful that non-mainstream voices are seldom represented in tech policy discussions, we have developed a formal process of refining our analysis with diversity panels, i.e., panels of individuals whose experience or expertise lies outside of the tech mainstream.

We engage with law and policy researchers to uncover assumptions jurists and policymakers might hold that no longer make sense in light of the new technology.

We offer a set of conditional recommendations that depend upon the particular values and goals policymakers, broadly construed, are trying to achieve.

ONE: TOWARD A WORKING DEFINITION OF AR

Augmented reality is shaping up to be an important and widespread technology. Some specific examples of AR being marketed or developed today include: Google Glass, Microsoft's HoloLens, Sony's Smart EyeGlass, Meta's Space Glasses, Magic Leap, Navdy Automotive, Across Air, and Word Lens.



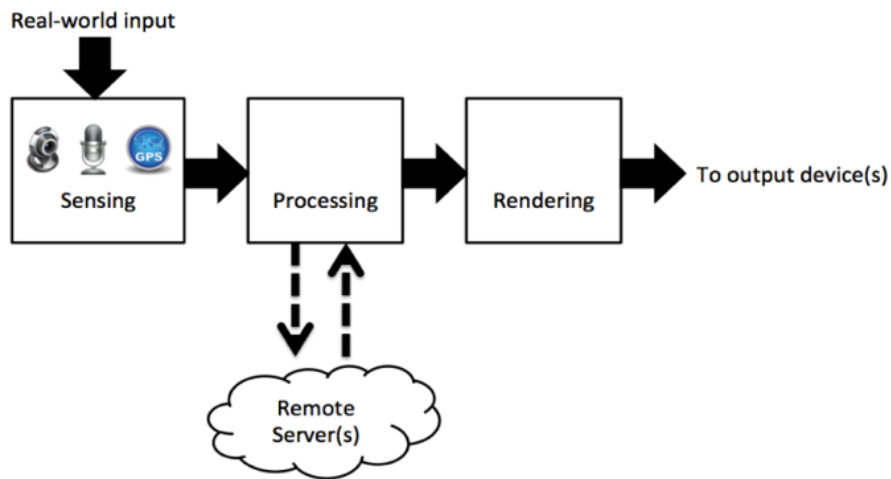
Google Glass, a heads-up augmented reality display.

(Image sources: left image from Google Glass invitation email, right image from <https://www.youtube.com/watch?v=d-y3bEjEVV8>)

As these examples show, there is no easy definition of “augmented reality”; AR is best understood as a class or family of technologies that tend to have certain common and distinguishing features. We have identified six such features, most of which are present in most AR systems:

1. *Sense properties about the real world.* The system will collect various forms of data about the world as the user experiences it. Sensors may include video (e.g., depth cameras, cameras worn on the body), audio, haptic input (i.e., detecting physical touch), location (e.g., GPS, GSM triangulation), motion, or wireless signals (e.g., WiFi, Bluetooth).
2. *Process in real time.* Inputs from the sensors will be analyzed and used by the system in real time. Some information may be stored for later analysis or sharing (e.g., life-logging), but at least some of the data is used in real time.
3. *Output (overlay) information to the user.* Information gathered and processed by the system will generally be overlaid on the user's usual perception of the world; this is unlike virtual reality, which entirely replaces the user's setting with a new environment. In augmented reality, information may be conveyed to the user via a variety of devices, including a screen, a speaker, or haptic feedback (e.g., vibrations, air pulses). Researchers are even experimenting with visual feedback via contact lenses.
4. *Provide contextual information.* The information provided by the system to the user is contextual and timely, meaning it will relate to what the user is currently experiencing. For example: real-time in situ language translation, ratings for restaurants passed on the street, or arrival time updates while waiting at the bus stop.

5. *Recognize and track real-world objects.* The feedback will tend to track or process real-world objects or people in the user's view. For example, a facial recognition application may recognize faces and label them with names as the identified person moves through the user's field of view.
6. *Be mobile or wearable.* In the long term, we expect that many augmented reality systems will be wearable (e.g., AR glasses), and the majority of our analysis will focus on such systems. However, a system does not need to be wearable to technically be considered an AR system; mobile options include some smartphone applications and heads-up displays in cars. Similarly, the Xbox Kinect facilitates many AR applications, but is not itself mobile.



Overview of an augmented reality application or system
 (Image source: Augmented Reality: Hard Problems of Law and Policy)

This definition helps distinguish AR from previous and constituent technologies while encompassing a burgeoning variety of AR applications. AR is still a young technology, so design choices are evolving and have yet to become industry standards. These choices in turn will drive the capabilities of AR, its impact on people, and its ramifications for law and policy.

One design parameter that warrants special attention is the *operating system* (OS), i.e., the underlying software platform or firmware that dictates the system's basic architecture. Full-fledged AR will run on an OS—Google Glass, for instance, runs on a modified version of the Android OS and Microsoft's HoloLens will run on Windows 10.

The configuration of the OS drives other choices: whether to open the platform to applications (apps) or other third party innovation; whether and how long to store information; whether to process information locally on the device or on a remote server; and whether the information is processed by a computer system, by a person (through crowdsourcing), or some combination. We also note that data collected from an AR system can be combined with other data sets, and brought together in the processing and display.

Further design decisions relate to the *interface*, how the user perceives and interacts with the system. Such issues include: whether and how to notify the user—and others—that an AR system is actively recording; the user's control over various parameters such as the access the system gets to her social media, email, or other accounts; and what information she perceives. Each of these choices affects not only the performance and versatility of the system but the ways AR interacts with the contemporary legal system.

TWO: HOW AR AFFECTS HUMAN EXPERIENCE



BMW Augmented Reality prototype guides a user through a repair procedure.

(Image source: BMW Augmented Reality, www.youtube.com/watch?v=P9KPJIA5yds)

AR systems change how people can interact with each other and their environment. AR users perceive more (or less) than the person, place, or object actually before them, as information is overlaid on their view of the world. They can also record their surroundings for future analysis. AR can confer new information in real time or alter the user's actual experience or skillset: someone who is poor at remembering faces can receive a prompt reminding him of where he met the person he is talking to, while someone without technical knowledge of cars could be guided through how to change a fan belt.

AR also changes the experiences of the people around the user, whose features and actions may now be recorded and analyzed—with or without notice, depending on designer choices. Moreover, AR makes it possible that two or more people perceive the same environment differently. One person may perceive an environment in an augmented state, while another may not. Two people may both experience an augmented space, but their versions may consist of different information overlays.

Not everyone experiences AR the same way, i.e., as “augmenting” reality by introducing new sensory information. For some populations—notably, those living with disabilities—AR may fully or partially replace a sense. Thus, for instance, an assistive technology may vibrate as people or objects approach or convert auditory information to visual stimuli. Those living with disability may come to rely upon these signals, such that their sudden interruption could create an inconvenient or even dangerous sensory deficit. These and other non-mainstream experiences must be kept firmly in mind when enumerating the potential use cases of AR, and as we contemplate rules and possible exceptions.

AR systems may also prove both empowering and disabling for a given population depending on the context. For example, AR could empower incarcerated youth providing a wider range of educational experiences, including hands-on work that would otherwise require intense investment in physical tools or spaces. But AR could also hinder these populations to the extent their arrest or incarceration records are rendered more visible to friends and neighbors, landlords, law enforcement, or prospective employers. Similarly, though AR could reveal, distract, and imperil people in new ways, it could also empower them by permitting them to record their surroundings, communicate with others, and gain information while keeping alert and not looking down.

THREE: CHALLENGES FOR LAW AND POLICY

AR systems change human experience and, consequently, stand to challenge certain assumptions of law and policy. The issues AR systems raise may be divided into roughly two categories. The first is collection, referring to the capacity of AR devices to record, or at least register, the people and places around the user. Collection raises obvious issues of privacy but also less obvious issues of free speech and accountability. The second rough category is display, referring to the capacity of AR to overlay information over people and places in something like real-time. Display raises a variety of complex issues ranging from possible tort liability should the introduction or withdrawal of information lead to injury, to issues surrounding employment discrimination or racial profiling. Policymakers and stakeholders interested in AR should consider what these issues mean for them.


Issues related to the collection of information include:

1. *Reasonable expectations of privacy.* Constitutional, tort, and much statutory privacy law (such as wiretap restrictions) turn on whether an action violates a citizen or consumer expectation of privacy that society is prepared to accept as reasonable. The introduction of always-on recording devices into public and private spaces may cause societal expectations to shift in ways that further diminish privacy recourse. Or, alternatively, such devices may place so much pressure on the reasonable expectation doctrine that courts may choose to revisit its utility.
2. *The third party doctrine.* U.S. case law holds that citizens have a diminished expectation of privacy in information (e.g., bank records) that they willingly convey to a third party (e.g., the bank). AR systems have the potential to relay virtually all the activity of a user to third-party servers for storage and processing. As with the reasonable expectation of privacy test generally, this development will make government access to the daily lives of citizens easier or call the already scrutinized third party doctrine into further question.
3. *Free speech.* The First Amendment right to free speech is a well-known right backed by a highly developed body of case law; more and more courts are beginning to recognize a corollary right to record and gather information as a prerequisite to expression or as expressive conduct in itself. The right tends to be a function of two factors: where the recorder is, and who (or what) is the subject of the recording. The routine use of AR may, again, strain this burgeoning doctrine. There is also an attendant concern that government will take countermeasures and seek to block the use of AR, limiting its potential as a tool of accountability.
4. *Intellectual property.* The always-on recording of everyday life will inevitably capture work that is protected by copyright, trademark, or other intellectual property laws. This risk may be particularly acute in movie theatres, concerts, or business settings. The mere collection and analysis of protected work is unlikely to trigger liability for most AR functions—for example, an application that displays the name of the song the AR user is hearing. Yet we can imagine a diverse array of situations in which copyrighted work is broadcast or superimposed, or in which trade secrets are compromised. The fair use doctrine may be frequently implicated in some of these scenarios, but since fair use is considered a defense to claims of infringement—rather than an exemption—its applicability will always be tested only after the fact, complicating the decision-making of technologists and policy makers concerned with how to regulate AR use and functionality in real life contexts.

Issues related to display of information include:

1. *Negligence.* AR systems overlay information onto the world in real-time. If users rely on this information in error, or if the information distracts a user, then any resulting injury could lead to a cause of action for negligence. Imagine, for instance, a heads-up display in the windshield of a vehicle that places an advertisement over a stop sign. The defendant could be the user herself, the manufacture of the AR system, an individual app developer, or all three.
2. *Product liability.* As a general matter, information, ideas, or expressions do not qualify as “products” for the purpose of product liability. Thus, for instance, the publisher of a book about edible mushrooms is not liable to readers who eat poisonous mushrooms in reliance on the book. But AR systems do more than relay information; they blend information with everyday activities in ways that can blur the distinction between real and perceived environments and risk physical harm. This could lead to a new category of product liability at the intersection of information and object.

3. *Digital assault.* AR can make objects appear that are not there, and disappear objects that are. Accordingly, there is the prospect of purposely harming or instilling fear in an AR user. This photo shows a prank involving the superimposition of a spider onto someone’s hand—a mild example, but AR advertisers and others may decide to employ shocking visuals and other effects to get attention. Under current theories of American tort law, purposely placing someone in imminent apprehension of physical harm—even where no harm is possible—can constitute a tort.



(Image source: www.youtube.com/watch?v=WjWzDOZmmxU)

4. *Discrimination.* AR will make it possible for users to look up information about people and places in real-time. The information gained in this fashion could lead to adverse decisions that are normatively unfortunate and even illegal. Anti-discrimination laws prevent decision makers from factoring some pieces of information into their decisions—for example, in employment, housing, or a number of other contexts. Imagine a jurisdiction that does not allow employers to consider an applicant’s arrest record or marital status while making hiring decisions. An AR system that provided an applicant’s rap sheet or dating profile automatically would be problematic under this regime. In other cases, though a form of discrimination may not specifically be proscribed, it is often best to remain ignorant of such information to limit liability exposure. More broadly, the display of crime statistics, housing prices, or other information about a building or neighborhood could further isolate urban and other environments by conveying an often false sense of danger.

FOUR: (CONDITIONAL) RECOMMENDATIONS

This whitepaper defines AR, provides a basic technical description, envisions implications for human experience, and identifies some of the unique or acute legal and policy issues raised by the possibilities of AR. This final section discusses options available to various policymakers who have an interest in promoting or regulating the technology. We identify a set of best practices and explain the interaction between various technical decisions and the law and policy landscape. The recommendations are “conditional” in the sense that they do not purport to advance any particular vision, but rather provide guidance that can be used to inform the policymaking process, regardless of the specific values any individual policymaker—or group of policymakers—may seek to advance.

1. *Build dynamic systems.* AR technology is advancing rapidly. Today’s systems should be flexible and capable of updating in the face of technical and cultural change. Law and policy itself, to stay relevant, should not assume a fixed instantiation of AR for all time.
2. *Conduct threat modeling.* Adversaries succeed in defeating systems precisely by finding behaviors that designers didn’t anticipate. A careful and thorough model of who might seek to compromise AR systems and how—without preconceptions (“no one would ever do that”)—is critical to managing privacy and security risks. Threat modeling is especially crucial where, as with AR, a system compromise could cause physical harm.
3. *Coordinate with designers.* Neither the design of AR nor the design of technology policy should occur in isolation. Technologists may not be aware of certain values held by policymakers and may not appreciate the legal import of a particular design decision. Policymakers, in turn, need an accurate mental model of the technology in order to make wise decisions—including omitting to take action.
4. *Consult with diverse stakeholders.* People will experience AR very differently, depending on their characteristics, experiences, and capabilities. Academia and industry interested in widely useful AR should expressly consult with diverse populations and solicit and incorporate their feedback.
5. *Acknowledge tradeoffs.* Design decisions matter. A system that is open to third party analysis or contribution—from open source code to an app store—may promote greater freedom and innovation, even as it opens consumers up to potentially malicious applications. Long-term information storage, cloud processing, and other advanced data processes may result in faster performance or more complex functionality, but at potential costs to privacy and free speech. Perfection can be the enemy of the good, but identifying important values and describing how architecture affects them is critical to responsible innovation in and beyond AR.

CONTRIBUTORS

The following people from the Tech Policy Lab contributed to this whitepaper:

Ryan Calo, JD

Tamara Denning, PhD (computer science)

Batya Friedman, PhD (information science)

Tadayoshi Kohno, PhD (computer science)

Lassana Magassa, PhD candidate (information science)

Emily McReynolds, JD, LLM

Bryce Clayton Newell, JD, PhD (information science)

Franziska Roesner, PhD (computer science)

Jesse Woo, JD

Thank you to the participants in our Diversity Panels, whose insights contributed greatly to this whitepaper.

Cover Image Credit: Leonard Low, via Wikimedia Commons, <http://bit.ly/1ii6Xtu>

To learn more about the Lab, visit us at techpolicylab.uw.edu.

BIBLIOGRAPHY

Books, Articles, & Papers

Babaguchi, N., Koshimizu, T., Umata, I., & Toriyama, T., *Psychological study for designing privacy-protected video surveillance system: PriSurv*, In: "Protecting Privacy in Video Surveillance" at 147 (Springer-Verlag 2009).

Brassil, J., *Technical challenges in location-aware video surveillance privacy*, In: "Protecting Privacy in Video Surveillance" at 91 (Springer-Verlag 2009).

Butler, D. J., Huang, J., Roesner, F., & Cakmak, M., *The privacy-utility tradeoff for remotely teleoperated robots*, In: Proceedings of the 10th Annual ACM/IEEE International Conference on Human-Robot Interaction (2015).

D'Antoni, L., Dunn, A., Jana, S., Kohno, T., Livshits, B., Molnar, D., Moshchuk, A., Ofek, E., Roesner, F., Saponas, S., Veanes, M., & Wang, H. J., *Operating system support for augmented reality applications*, In: Proceedings of the USENIX Workshop on Hot Topics in Operating Systems. (2013).

Deng, J., Krause, J., and Fei-Fei, L., *Fine-grained crowdsourcing for fine-grained recognition*, In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2013).

Denning, T., Dehlawi, Z., & Kohno, T., *In situ with bystanders of augmented reality glasses: Perspectives on recording and privacy-mediating technologies*, In: Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems (2014).

- Felt, A. P., Ha, E., Egelman, S., Haney, A., Chin, E., & Wagner, D., *Android permissions: User attention, comprehension, and behavior*, In: Proceedings of the Symposium on Usable Privacy and Security (2012).
- Friedman, B., Kahn Jr, P.H., & Borning, A., *Value Sensitive Design and information systems*, In: "Human-computer interaction in management information systems: Foundations" at 348 (Armonk 2006).
- Friedman, B., Smith, I. E., Kahn Jr, P.H., Consolvo, S., & Selawski, J., *Development of a privacy addendum for open source licenses: Value Sensitive Design in industry*, In: Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (2006).
- Gupta, S., Morris, D., Patel S.N., Tan, D., *Airwave: non-contact haptic feedback using air vortex rings*, In: Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (2013).
- Halderman, J. A., Waters, B., & Felten, E. W., *Privacy Management for Portable Recording Devices*, In: Proceedings of the Workshop on Privacy in Electronic Society (2004).
- Hayes, G. R., & Truong, K. N., *Selective Archiving: A Model for Privacy Sensitive Capture and Access Technologies*, In: "Protecting Privacy in Video Surveillance" at 165 (Springer-Verlag 2009).
- Hoyle, R., Templeman, R., Armes, S., Anthony, D., Crandall, D., & Kapadia, A., *Privacy behaviors of lifeloggers using wearable cameras*, In: Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (2014).
- Jana, S., Molnar, D., Moshchuk, A., Dunn, A., Livshits, B., Want, H. J., and Ofek, E., *Enabling fine-grained permissions for augmented reality applications with recognizers*, In: Proceedings of the USENIX Security Symposium (2013).
- Jana, S., Narayanan, A., & Shmatikov, V., *A Scanner Darkly: Protecting user privacy from perceptual applications*, In: Proceedings of the IEEE Symposium on Security and Privacy (2013).
- Jung, J., & Philipose, M., *Courteous Glass*, In: Proceedings of the Workshop on Usable Privacy & Security for Wearable and Domestic Ubiquitous Devices (2014).
- Kooper, R., & MacIntyre, B., *Browsing the Real-World Wide Web: Maintaining Awareness of Virtual Information in an AR Information Space*, 16 International Journal of Human-Computer Interaction 425 (2003).
- Mark Hornyack, P., Han, S., Jung, J., Schechter, S., and Wetherall, D., *"These aren't the droids you're looking for": Retrofitting Android to protect data from imperious applications*, In: Proceedings of the 18th ACM Conference on Computer and Communications Security (2011).
- McPherson, R., Jana, S., & Shmatikov, V., *No Escape From Reality: Security and Privacy of Augmented Reality Browsers*, In: Proceedings of the 24th International World Wide Web Conference (2015).
- Ng-Thow-Hing, V., Bark, K., Beckwith, L., Tran, C., Bhandari, R., & Sridhar, S., *User-centered perspectives for automotive augmented reality*, In: Proceedings of the IEEE International Symposium on Mixed and Augmented Reality (2013).
- Paruchuri, J. K., Cheung, S.-C. S., & Hail, M. W., *Video data hiding for managing privacy information in surveillance systems*, EURASIP Journal on Information Security (2009).
- Patel, S. N., Summet, J. W., & Truong, K. N., *BlindSpot: Creating capture-resistant spaces*, In: "Protecting Privacy in Video Surveillance" at 185 (Springer-Verlag 2009).
- Raval, N., Srivastava, A., Lebeck, K., Cox, L. P., & Machanavajjhala, A., *MarkIt: Privacy Markers for Protecting Visual Secrets*, In: Proceedings of the Workshop on Usable Privacy & Security for Wearable and Domestic Ubiquitous Devices (2014).

- Roesner, F, Denning, T., Newell, B.C., Kohno, T., Calo, R., *Augmented Reality: Hard Problems of Law and Policy*, In: Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (2014).
- Roesner, F, Molnar, D., Moshchuk, A., Kohno, T., & Wang, H. J. *World-Driven Access Control for Continuous Sensing Applications*, In: Proceedings of the ACM Conference on Computer and Communications Security (2014).
- Schiff, J., Meingast, M., Mulligan, D. K., Sastry, S., & Goldberg, K., *Respectful Cameras: Detecting Visual Markers in Real-Time to Address Privacy Concerns*, In: "Protecting Privacy in Video Surveillance" at 65 (Springer-Verlag 2009).
- Sutherland, I. E., *A head-mounted three-dimensional display*, In: Proceedings of the Fall Joint Computer Conference, American Federation of Information Processing Societies (1968).
- Templeman, R., Korayem, M., Crandall, D., & Kapadia, A., *PlaceAvoider: Steering first-person cameras away from sensitive spaces*, In: Proceedings of the 21st Annual Network and Distributed System Security Symposium (2014).
- Vilk, J., Molnar, D., Ofek, E., Rossbach, C., Livshits, B., Moshchuk, A., Wang, H. J., & Gal, R., *SurroundWeb: Mitigating Privacy Concerns in a 3D Web Browser*, In: Proceedings of the IEEE Symposium on Security and Privacy (2015).

Cases

- Chicago Lock Co. v. Fanberg*, 676 F.2d 400 (9th Cir. 1982)
- Glik v. Cunniffe*, 655 F.3d 78 (1st Cir. 2011)
- Katz v. US*, 389 U.S. 347 (1967)
- Smith v. Maryland*, 442 U.S. 735 (1979)
- United States v. Jones*, 132 S.Ct. 945 (2012)
- Winter v. G.P. Putnam's Sons*, 938 F.3d 1033 (9th Cir. 1991)