



2015

Augmented Reality All Around Us: Power and Perception at a Crossroads

Marty J. Wolf
Bemidji State University

Frances Grodzinsky
Sacred Heart University, grodzinskyf@sacredheart.edu

Follow this and additional works at: http://digitalcommons.sacredheart.edu/computersci_fac

 Part of the [Computer Sciences Commons](#)

Recommended Citation

3. 2015: Wolf, M.J., Grodzinsky, F.S., Miller K.W. Augmented Reality All Around Us: Power and Perception at a Crossroads, *ACM Computers and Society*, Vol 45 issue 3, September, 2015. P. 126-131. Doi> 10.1145/2874239.2874257. Retrieved from <http://dl.acm.org/citation.cfm?id=2874257&CFID=610593124&CFTOKEN=62546977>

This Article is brought to you for free and open access by the Computer Science & Information Technology at DigitalCommons@SHU. It has been accepted for inclusion in Computer Science & Information Technology Faculty Publications by an authorized administrator of DigitalCommons@SHU. For more information, please contact ferribyp@sacredheart.edu.

Augmented Reality All Around Us: Power and Perception at a Crossroads

Marty J. Wolf
Bemidji State University
Bemidji, MN USA 56601
mjwolf@bemidjistate.edu

Frances Grodzinsky
Sacred Heart University
Fairfield, CT USA 06825
grodzinskyf@sacredheart.edu

Keith Miller
University of Missouri - St. Louis
St. Louis, MO USA 563121
millerkei@umsl.edu

ABSTRACT

In this paper we continue to explore the ethics and social impact of augmented visual field devices (AVFDs). Recently, Microsoft announced the pending release of HoloLens, and Magic Leap filed a patent application for technology that will project light directly onto the wearer's retina. Here we explore the notion of deception in relation to the impact these devices have on developers, users, and non-users as they interact via these devices. These sorts of interactions raise questions regarding autonomy and suggest a strong need for informed consent protocols. We identify issues of ownership that arise due to the blending of physical and virtual space and important ways that these devices impact trust. Finally, we explore how these devices impact individual identity and thus raise the question of ownership of the space between an object and someone's eyes. We conclude that developers ought to take time to design and implement a natural and easy to use informed consent system with these devices.

Categories and Subject Descriptors

K.4.1 [Computers and Society]: Ethics.

General Terms

Human Factors.

Keywords

Augmented Reality, Augmented Visual Field Devices, Autonomy, Deception, Human Values, Identity, Informed Consent, Trust.

1. FRAMING THE DISCUSSION

This paper extends and elaborates an earlier paper, Grodzinsky, Miller and Wolf [1]. In that paper, we explored augmented visual field devices (AVFDs), using the following definition for *visual augmented reality* (AR): "...visual augmented reality involves projecting light in such a way that both natural light and artificial light enter the eye simultaneously, so that some objects seen in the visual field can be traced back to physical objects, and other objects seen are virtual objects, for which no physical object is the source of reflected light." Since that time, Microsoft announced the pending release of HoloLens, "the world's most advanced holographic computing platform" [2]. HoloLens seemingly will project holographic images into the physical space that are visible

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

to at least the wearer of the HoloLens. In addition Magic Leap has recently filed a patent application for technology that rather than have the user viewing artificial light emanating from a screen, the device will project light directly onto the wearer's retina [3]. The holy grail with all of these technologies is to create an environment where the user interacts with virtual and physical objects in a natural, seamless way. It appears the goal of many of these technologies is to make the virtual objects as similar to the physical objects in the immediate environment, to the point that the user is unable to distinguish between the virtual and the real in his or her interactions.

Certainly the cameras that are incorporated into AVFDs are an obvious point of concern regarding the technology. Denning, Dehlawi, and Kohno [4] conducted a small experiment of reactions bystanders have to cameras and recording devices. Their work revealed that the newness and unfamiliarity of these devices caused bystanders to view them differently from other recording devices such as mobile phones. Their mock recordings took place in a cafe and many bystanders thought the researchers ought to be required to get permission before recording. Some bystanders showed an interest in a (hypothetical) device that would block such recording.

While Denning et al. focused solely on recording components, AVFDs certainly will contain other familiar components such as GPS. Some of the ethical concerns we raise are not new to AVFDs; however, the nature of these concerns change when these technologies are combined into a single device with proposed components of AVFDs such as holographic projectors and retina projectors. Often times promoters of these technologies speak of the advantages the individual user of the device will experience. There seems to be little analysis of both the potential disadvantages to the individual user and almost no analysis to the impact these devices might have in larger groups and on social structures. We address some of these ethical concerns here.

Friedman and Kahn [5] examined augmented reality using seven human values they predicted would be important for understanding the ethical import of AR. In our earlier paper, we explored three of those seven: psychological well-being, physical well-being, and privacy. In this paper, we will concentrate on the remaining four values in Friedman and Kahn's list: deception, trust, informed consent, and ownership. We will also draw from the additional values that Friedman and Kahn suggested in a subsequent paper [6]. They include freedom from bias, universal usability, autonomy, identity, calmness, courtesy and environmental sustainability. AVFDs (especially future devices) will embed several of these values.

Our discussion focuses on four groups of stakeholders involved with AVFDs: developers, users (both individually and collectively), non-users who are in sight of users, and society as a whole. We will use the term "developers" in a broad sense, meant

to include at least designers, software engineers, and managers of the companies making these devices. We will call non-users who can be physically (not virtually) seen by AVFD users at “the watched.”

2. DECEPTION

Deception is like the tango – it takes at least two, a deceiver and a deceived. In [7] we considered deception to be “an intentional, successful attempt by developers to deceive users, and a misapprehension by people other than the developers.” Consistent with Lynch [8], deception requires a misleading act that is “*willful or non-accidental*. So, X deceives Y with regard to f only if X willfully causes Y to fail to believe what is true with regard to f.” It is important to note that deception is not inherently bad. As we noted in [7], developers regularly hide implementation details from users to make the user experience more familiar (e.g. the use of a the folder and file metaphor for the file system). We called this a benign deception. In this section we describe three possible deception relationships that we think are both likely and ethically significant with AVFDs. All three relationships involve users, but one of them includes developers, and one of them includes the watched. AVFDs seem to strain Lynch’s requirement for deception that the act be willful or non-accidental in that a user of an AVFD may make a willful act of using one, but have his/her reality become so intertwined with a virtual component that the possibility of willfully deceiving someone is no longer a conscious choice.

A more interesting take on deception for this application might be that of Mark Wrathall [9]. Wrathall offers insight into deception as a perceptual experience. Wrathall writes, “In the genuine perceptual experience, the phenomenal character of things corresponds to the way things actually are. One then accounts for deceptions by treating them as the presentation of a certain phenomenal character in the absence of the objects necessary to make that presentation true” [9]. He goes on to explain that “when we are deceived, it’s because the thing really looks like what we take it as.” So deceptions, in this sense, have to do with misperceptions. It is how we view the world and how the world is presented to us [7]. This raises an interesting question for the case of AVFDs where a genuine perceptual experience includes not only phenomenal character of things but also the virtual. Everyone’s perception of the same object may be different because of what is virtually added. We would not call this a misperception but rather an augmented one. So, how can we tell if an augmented perception is a deception? In a certain sense, AR is all about fooling the user’s eyes and brain. So where do we cross the line? Great care should be taken to help users be discerning consumers of this new information.

2.1 Developers May Deceive Users

AVFD developers have several kinds of power over users. First, the developers know many technical details about the devices, details that are not obvious to most users. Because of this information (and power) imbalance, developers could deceive users about the capabilities and sophistication of the AVFD and its algorithms. This kind of deception would not be distinctive to AVFDs, but is common to all high tech devices. However, the nature of AVFDs, the intimacy of changing what people see, might increase the ethical significance of this particular technology deception.

The ancient slogan “seeing is believing” [10] illustrates another way that AVFD developers might deceive users. Should

developers succeed in engineering the AR experience in such a way that augmented reality is indistinguishable (or nearly so) from physical reality, users might be deceived into believing in the physical existence of what they see, even though it is not physically present. In the case of devices that display light directly on the user’s retina, the intention to deceive cannot be eliminated from the nature of the AVFD. The user cannot distinguish the two different sources of light. It will take other cues for the user to determine the virtual from the physical.

Regardless of whether a virtual object is a holographic image or being displayed directly onto the user’s retina, the developer takes on additional responsibility for the veracity of any information attached to the object. Either purposefully, or carelessly, developers could deliver bogus information to users. It may very well be that users who see that information called up instantly and effortlessly into their visual space will be inclined to give that information the benefit of any doubts about the information’s accuracy. One way to mitigate this concern would be to make it obvious to the user the nature of control that she has over information. Yet one of the developmental difficulties is determining a convenient way for a user to provide input into an AVFD. Shortcomings in this feature lead to more control for the developer and less for the user. Therefore great care should be taken to help users be discerning consumers of the information they are perceiving in order to mitigate the potential for developers to routinely deceive the users.

2.2 Users May Deceive the Watched

In considering how AVFD developers may deceive users, we concentrated on AR outputs to the user’s eyes. In considering how AVFD users may deceive the watched, we also consider AR visual inputs, real time video taken from the user’s viewpoint. The potential for privacy invasion was one of the reasons Google Glass users were not universally welcomed into public spaces [4, 11]. Users, recording members of the watched, might explicitly or implicitly lie about their actions or intentions.

In addition, users might misrepresent what they are seeing via their AVFD. We can envision many scenarios in which a user either has or might have information that non-users do not have. A user might be asked about that information, or a user might volunteer it. Either way, the user might misrepresent the presence or absence of the requested information, or misrepresent the information. “Yes, I can see that...” could be used as a method of establishing authority and seeking the power of information (whether the information is true or false). Rather than create an atmosphere of trust, these potentials for deception create one of distrust and uneasiness.

2.3 Users May Deceive (Other) Users

One of the interesting aspects of AVFD systems is the potential for multiple users (who will probably have to be using similar, if not identical, systems) to interact in the virtual space overlaid on their individual physical views. So, for example, we are told that we will be able to play virtual chess, or laser tag, with each other. But it does not take a great deal of imagination to anticipate that some AVFD users who share virtual space with other users could rig the common virtual experience to their individual advantage. For example, one laser tag participant may find a way to have the game unfairly slanted to his or her advantage. There are numerous ways a virtual poker game could be used to cheat opponents. Users might purposefully share inaccurate information (for

example, information about other people in the room) that would be displayed on nearby AVFDs.

A single user of an AVFD may choose to deceive him/herself. Someone may choose to adorn him/herself with opulent jewels or even keep a long dead pet close at hand. As virtual worlds and the physical world become increasingly blended, questions about what is real will begin to change. Each of us can create our own blend of physical and virtual to create our own realities. Underlying assumptions about all of us sharing the same reality will no longer hold.

3. INFORMED CONSENT AND AUTONOMY

Medical informed consent [12] implies knowledge of the intended intervention, awareness of possible risks and benefits, and an explicit declaration of agreement that the procedure go forward. Applying this idea to the use of AVFDs, several aspects already discussed seem relevant. Consistent with the sensibilities of the subjects of Denning et al.'s experiment [5], there is a case that both AVFD inputs and outputs should be considered for informed consent. First, a user who uses an AVFD to record images or audio should do so only with the consent of people included in the recording, particularly if the recording is going to be shared. This is further complicated when the recording includes holographic images that appear to be a real part of the physical space. A simple, uniform method for describing the nature of the recording, what is allowed and what is not allowed and the actual obtaining of consent from the watched, especially in large crowds, seems to be no simple feat.

Furthermore, if a developer or a user is responsible for changing another user's virtual space, it should be clear to the affected user that this change is taking place. Surely some such changes would be well known by the users involved; if a user bought a virtual chess program, and the developer delivered an appropriate set of virtual objects for the players, no formal informed consent would be necessary as it is implicit in the product. But if a developer or a user X controlled the virtual space in such a way that all watched individuals were scanned, and otherwise private information appeared in X's virtual view, then the watched individuals should be asked for their consent, or it should not be done. We can imagine scenarios (for example, in an emergency room) where watched patients might be willing to give such consent to medical staff with AVFDs. But we can also imagine some patients in an ER who would refuse consent. Either way, it should be an option, not a requirement, for treatment.

We can envision situations in which people would waive AVFD informed consent. For example, some AVFD enthusiasts might want to gather and experience each other's virtual manipulations. If fellow users were trusted, or if enthusiasts did not care about the consequences of giving up their control of virtual space, they might mutually agree to a common license (among themselves) for a wide-open experience. As long as such agreements are explicit and mutually agreed upon, we do not see an ethical problem.

We can also envision scenarios in which someone was coerced into using an AVFD. The coercion could be economic, where as part of your job requirements you had to agree to training with an AVFD. The coercion could be legal; for example, probation could be granted only if a prisoner was willing to undergo AVFD therapy, therapy that was designed to induce revulsion at certain triggering situations. In these types of cases, authorities

(commercial or governmental) may reason that the greater good (of a corporation or a polity) trumps the need for voluntary informed consent. We are suspicious of such reasoning, and we contend that great care should be taken when forcing AVFD experiences on to individuals.

While it is likely that in general people will not be forced to use AR devices, we can envision certain contexts where such use may be encouraged or even required, for example at work or in school. An AR device that can "pin" holographic objects in the real world and can allow users to interact with both virtual and physical objects simultaneously seems to offer a potentially valuable learning environment. A student in a class which requires interaction with a pinned hologram would seem to have little choice but to acquire and don an AVFD. Using such a device as a classroom tool is not necessarily ethically problematic if all students have access. However, "having access" may be more complicated than simply having a device to use; some students may not be able to benefit from an AR device. Blind students are an obvious example, but some sighted students might have adverse reactions to an AR device, including headaches or dizziness; how will such students be treated if an AR experience is a required part of a curriculum? Teachers have a tradition of guiding students' learning in similar ways. However, issues of autonomy do creep into this situation. We need mechanisms to determine the level of control each student should have. The teacher and the school will also exhibit some level of control over the experience, with one or the other potentially having complete control over each student's use of the device. As a collaborative and learning tool, it may be useful for students to see the interactions and the results of interactions that other students initiate.

4. OWNERSHIP

Several ownership issues arise surrounding AVFDs. First, will AVFDs be owned (like most computer hardware) or leased (like much proprietary software)? We assume that the AVFD hardware will be owned, but that much of the software will be leased. Proprietary software is likely not to be readily accessible for users or for the watched; therefore, there may be interest in having at least some AVFD software be free or open source software (FOSS). We will not reprise the arguments for and against proprietary and FOSS solutions here, but this is a venue where those arguments will again play out, affecting the balance of power between developers and users, and to a lesser extent between users and the watched.

In previous sections we pointed out the possibilities for deception and informed consent situations having to do with AVFD users recording images and sounds from the watched and from other users. This aspect of AVFDs can be viewed as an ownership issue: who owns my recorded image and voice? Legally, particular instances of this argument may turn on where the AVFD is deployed. If the recording takes place in a public space, then the watched may not have a presumption of privacy; if the recording takes place in a space that is not legally designated as public, then there may be a presumption of privacy. However, we suspect that an ethical analysis would be more restrictive of a user's "right" to the use of the watched's images and voices. For a more complete discussion of AVFDs and privacy, see [1].

The issues of ownership of devices and recorded images for AVFDs are interesting, but closely related to issues with previous devices. Graham, Zook, and Boulton [13] demonstrate the power that comes with one augmented reality technology, Google Maps,

by demonstrating how Google shows and describes places differently depending on the language one uses to view a particular place. A more distinctive ownership issue for AVFDs is: who has legitimate claims to the virtual space (what the users see)? We assume that a user should have some claim to that space, since it is his/her device, and since his/her eyes and visual cortex are most immediately impacted by the virtual image. However, the developers of the device work to design and deliver that virtual environment, and they might also make a claim of ownership; the developers clearly do have control, especially initially, on that virtual space. If some AVFD applications require real time Internet sharing (similar to what gaming systems use for multiplayer games), again that virtual space is claimed by both developers and users.

This sort of sharing also suggests a need for open standards. Proactive work on how virtual objects and experiences are to be represented and shared will allow for users with different brands of devices to be unencumbered by those differences. There is a need in the AVFD arena for the same sort of frictionless interaction that we experience while texting, making phone calls and sharing photos.

In cases where both developers and users may have possibly legitimate claims to ownership, we think it is vital for the participants to have explicit agreements about the ownership of the virtual space. It may be that in particular applications (such as shared AVFD games), users will be content to relinquish control in order to enter into a group experience. In other applications (for example, a surgeon using AVFD during an operation), users may demand a much higher degree of control, especially when they are responsible for critical decisions based partly on information delivered by an AVFD. In both these cases, the stakeholders can act ethically, but only when the agreements are explicit, appropriately detailed, and understood by all parties.

One virtual space of particular interest is that surrounding existing physical objects. The Artvertiser project started by Julian Oliver [14] seeks to “improve reality” by placing virtual art over advertising in public spaces through the use of AVFDs. While the virtual art is visible only to the wearer of the AVFD, it does “prevent” the wearer from seeing the advertisement on the billboard. An advertiser might argue to the AVFD developer that such an ability ought to be blocked on the AVFD. Since so much software on portable devices is largely supported by advertising, this sort of feature might lead to a decrease in economic support of software available for AVFDs or an increase in the price of that software. On the other hand, there is no clear argument that one ought to be subjected to advertising in public spaces. Even without AVFDs, people can avert their eyes. Yet, the intriguing question remains, should someone be allowed to own the visual experience in a public place?

Closely related to that question is perhaps the most important aspect of AVFD ownership--that of an individual's ownership of his/her own perception. In some sense, donning an AVFD allows someone (or something) to radically alter what the individual perceives. This temporary surrender of control has analogs in other technology. When we see a film at a theatre, when we watch television, and when we listen to an iPod, we are giving control over one or more of our senses to a machine and the sociotechnical system of which that machine is a part. But the distinctive mixture of physical and virtual that is delivered by AVFDs may be seen as a qualitatively greater surrender. And if it becomes commonplace to make that surrender on a daily, or even continuous basis, then part of who we are, and much of what we

see, will be “owned” outside of ourselves. That is a major ethical issue with power at its core..

5. TRUST

AVFDs are artifacts that mediate our perception of reality. According to our Object Oriented model of Trust [15], they would fall under the category of human to human trust mediated by electronic means. There we state: “The people who design, develop, or deploy a computing artifact are morally responsible for that artifact, and for the foreseeable effects of that artifact. This responsibility is shared with other people who design, develop, deploy or knowingly use the artifact as part of a sociotechnical system.” [15] What is the impact on trust?

There are two trust relationships that must be considered: trust between users and developers; and trust between users and other individuals (some of whom may be users themselves, and other individuals who are not users). Both the developers and users must take on moral responsibility for the artifact. That is, developers of AVFDs should have as an accepted goal: examination of the effects of that artifact on society and performance of their functions with the appropriate standard of care. A subgoal here would be transparency: developers being honest with others about the capabilities of the device. Users who trust developers will buy their products and use them with confidence. However, if the user performs certain actions based on the trust he/she has in the artifact, and if that trust is misplaced (i.e., the developer is manipulating the end-user and does not have the user's best interests at heart), then there is a violation of trust [8]. In the second trust relationship, individuals must trust that users in public are employing the device in an ethically acceptable way.

Another issue of trust involves epistemic trust. How do we know what we know from our perceptions through AVFDs? Can we trust what we perceive to be true? Judith Simon says that “trust and knowledge are fundamentally entangled in our epistemic practices. Yet despite this fundamental entanglement, we do not trust blindly. Instead we make use of knowledge to rationally place or withdraw trust. We use knowledge about the sources of epistemic content as well as general background knowledge to assess epistemic claims. Hence, although we may have a default to trust, we remain and should remain epistemically vigilant; we look out and need to look out for signs of insincerity and dishonesty in our attempts to know” [16]. This statement could apply to the user's relationship with the developer. It is more difficult to trust what we see as true when the virtual and real are entangled and our world is mediated through a device. How does what we know impact what we perceive and conversely how does what we perceive impact what we know? The answer to these questions will affect whether we trust what we see through the AVFD.

6. IDENTITY

In addressing issues of identity, we note that AR devices may help individuals establish their own identities. There is the potential for a deep blending of the physical and virtual self. In the physical world, people use jewelry, body piercings, tattoos, and ear lobe gauging to distinguish themselves and establish at least part of their identity. People use posts on Pinterest, FaceBook, Twitter and Instagram to create a virtual part of who they are as individuals. AR devices open up the possibility of pinning these sorts of identity-creating virtual items to one's physical self, so that anyone with a compatible AVFD will see the pinned objects

when you are viewed. It could become similar to having “virtual jewelry,” with vendors competing to offer increasingly sophisticated decoration that shares all of the properties of information such as being easily and quickly updated. Those viewing someone who is virtually decorated through an AVFD will see that person as part physical, part holographic and perhaps be unable to distinguish between the two.

Of course, that is the ideal. A person ought to have autonomy over her/his identity, yet the AVFD through which the person is being viewed may be owned by someone else. At the very least there is the opportunity for the owner of the AVFD to use a different virtual accoutrements on a person than those selected by the watched. The possibility of decorating *others*, especially without their consent, seems fraught with difficulties and potential abuse. The potential problem becomes even more pronounced in a group setting. This seems to be especially true in a setting, such as a school, where bullying takes place. This technology opens up new avenues for cyber-bullying.

Ethical principles surrounding identity seem to collide with ethical principles surrounding public spaces. In the case of a public space there is a reasonable argument that an individual can choose to use an AVFD to obscure or replace an advertisement in a public space. In some sense the person lays claim to the visual space between the AVFD and up to, but not including the advertisement. On the other hand, when the AVFD user is viewing another person in a public space, the user’s right to control their own visual experience comes up against the watched’s identity and autonomy. As in so many questions about technology and people, power relationships are clearly important. Inasmuch as AVFDs empower individuals to thrive, there is a positive effect; inasmuch as AVFDs are used to enhance the power of those already powerful to the detriment of the less powerful, there is a negative effect.

7. ETHICAL CHOICES

In order to elaborate some of the ethical choices to be made with AVFDs, consider the following scenario: developers have set up a system that includes multiple users and the developers themselves to interact using interlinked AVFDs, sharing a physical and a virtual space together. Two examples of such a situation could be a developer, a surgeon, and a group of medical students inside an operating room; or a group of gamers and a developer in an outdoor setting playing a first person shooter game. What can we say about the actions of the developers in this situation?

First, the developer has at least two kinds of control in these situations: first, the developer controls the initial configuration of the system, including what the users will see (virtually), and how much control each user and the developers have over those virtual images. (In this paragraph, we will use “images,” but in many AVFDs, there could also be sounds added.) The second kind of control is real time, after the AVFDs are deployed, and the users are inhabiting the same physical and virtual space. In a move toward simplicity, the developer might decide that no one’s virtual images take precedence and thus block everyone’s. This option certainly detracts from the value of AVFDs. At the other extreme, the developer could allow everyone’s virtual images to be seen. This also seems to detract from the value of AVFDs as such an experience would be visually cluttered and noisy.

For a more nuanced look, consider the interests of a developer D and a user U (who is not a developer). Assume that D wants to associate certain virtual images V to U so that anyone looking at

U with one of the AVFDs in the system will see V (virtually) as well as U (physically). Several different situations arise:

1. U does not like some aspects of V, and objects to D, either before V is shown to others, or after V is shown to others. Whose preferences are likely to take priority? That probably depends on the situation, and on the power relationships outside the AVFD system. For example, in the medical situation, the surgeon will probably have a great deal to say about how s/he is presented to others, but a medical student might not have any say. In a gaming situation, users might have some latitude for some aspects (for example, they might make up a gaming ID that is virtually attached to them), but not for other aspects (for example, the game may insist on projecting their current life force). Deciding whose preferences *should* take priority, the ethical question, will be situation specific; however, we contend that developers should negotiate these kinds of issues early and often during development and deployment.
2. Now assume that two users, U1 and U2, have the power (granted by the developer) to change virtual images associated both with themselves and with each other. Again, conflicts can occur when one of the users “decorates” the other with images that the decorated user finds objectionable. We think it is central to the ethics of this situation what agreements the users entered into when they joined in the AVFD system. If they agreed to subject themselves to this decoration by others, then they probably may not have much to complain about. If U2 objects to U1’s decoration of U2, U2 can try to negotiate with U1 to remove or change the decoration, or U2 can withdraw from the system.
3. In both case 1 and case 2, the issues can be framed as informed consent. Thus a system that informs U2 of U1’s desires and allows U1 to either consent or not seems to be called for. This option has the positive of forcing an interaction. It does not seem to impinge excessively on U1. In the end, it opens the opportunity for collaboration, allowing both U1 and U2 to potentially thrive. This approach impinges on the developers, forcing them to design an entire system for this sort of exchange to take place. This is an interesting case for the ownership issue as well. Certainly, one ought to expect bullies and trolls to avoid this sort of thing, and it would be unreasonable to expect this sort of system to not be hacked. If the software were FOSS, then it would be easy for the bullies and trolls to avoid informed consent. Proprietary software, on the other hand, would make that more difficult.

Traditionally, the question of ownership of the space between an object and someone’s eyes has not been called into question. AVFDs have the potential to force us to consider that question. It opens up new opportunities for individual freedom for AVFD users (e.g., one can avoid being bombarded by advertisements in public spaces), and also potential hazards for the watched who could be seen not as they physically and virtually project themselves, but rather as the AVFD user would like. This is a collaborative environment of public and private, virtual and real. Developers, and the systems that they produce as part of AVFDs, will have an important role to play in the environment that surrounds these devices.

8. CONCLUSIONS

The phrase “I can’t believe my eyes” is meant to say that something is extraordinary, surprising, and unexpected. But if it becomes commonplace not to believe our eyes due to AR devices and policies that allow others to control what we see, we think that we will be engaging in a risky socio-technical experiment. We contend that such issues should be debated now, not after AR devices become commonplace.

9. REFERENCES

- [1] Grodzinsky, F. S., Miller, K. W., and Wolf, M. J. 2014. Augmented reality in your eye: Google Glass, Space Glasses, and beyond. In *CEPE 2014*.
- [2] Microsoft. 2015. Microsoft HoloLens. <http://www.microsoft.com/microsoft-hololens/en-us>.
- [3] Abovitz, R., Schowengerdt, B. T., and Watson, M. D. 2015. Planar waveguide apparatus with diffraction element(s) and system employing same. <http://www.faqs.org/patents/app/20150016777>.
- [4] Denning, T., Dehlawi, Z., and Kohno, T. 2014. In situ with bystanders of augmented reality glasses: perspectives on recording and privacy-mediating technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 2377-2386. DOI=10.1145/2556288.2557352 <http://doi.acm.org/10.1145/2556288.2557352>
- [5] Friedman, B., and Kahn, P. H., Jr. 2000. New directions: A value-sensitive design approach to augmented reality. In *Conference Proceedings of DARE 2000: Design of Augmented Reality Environments*. ACM, New York, NY, USA, 163-164.
- [6] Friedman, B., and Kahn, P. H., Jr. 2003. Human values, ethics, and design. In J. A. Jacko and A. Sears (Eds.), *The human-computer interaction handbook*. Mahwah, Lawrence Erlbaum Associates, NJ, USA, 1177-1201.
- [7] Grodzinsky, F. S., Miller, K. W., and Wolf, M. J. 2015. Developing automated deceptions and the impact on trust. *Philosophy and Technology*. 28, 1, (Mar. 2015), 91-105. DOI=10.1007/s13347-014-0158-7.
- [8] Lynch, M. P. 2009. Deception and the nature of truth. In M. Clancy (Ed.), *The Philosophy of Deception*. New York: Oxford University Press, 188-200.
- [9] Wrathall, Mark A. 2009. On the existential positivity of our ability to be deceived. In M. Clancy (Ed.), *The Philosophy of Deception*. New York: Oxford University Press, 67-81.
- [10] Ammer, C. 1997. Seeing is believing. *The American Heritage Dictionary of Idioms* Houghton Mifflin Harcourt Publishing Company.
- [11] Kelly, H. 2013. Google Glass users fight privacy fears. <http://www.cnn.com/2013/12/10/tech/mobile/negative-google-glass-reactions/index.html>.
- [12] Dictionary.com 2015. "informed consent," in *Dictionary.com Unabridged*. Random House, Inc. [http://dictionary.reference.com/browse/informed consent](http://dictionary.reference.com/browse/informed%20consent).
- [13] Graham, M., Zook, M., and Boulton, A. 2013. Augmented reality in urban places: contested content and the duplicity of code. *Transactions of the Institute of British Geographers*. 38, 3, 464-479.
- [14] The Artvertiser. <http://theartvertiser.com/>.
- [15] Grodzinsky, F. S., Miller, K., and Wolf, M. J. 2012. Moral responsibility for computing artefacts: ‘the rules’ and issues of trust. *Computers and Society*. 42, 2, (Dec. 2012), 15-25. DOI=10.1145/2422509.2422511
- [16] Simon, J. (2010), The entanglement of trust and knowledge on the web, *Ethics and Information Technology*. 12, 343-355.