Philosophy & Technology (2022) 35: 19
https://doi.org/10.1007/s13347-022-00496-5

RESEARCH ARTICLE

Augmented Reality, Augmented Epistemology, and the Real-World Web

Cody Turner[1]

Received: 22 March 2021 / Accepted: 10 January 2022 / Published online: 15 March 2022

Abstract
Augmented reality (AR) technologies function to 'augment' normal perception by superimposing virtual objects onto an agent's visual field. The philosophy of augmented reality is a small but growing subfield within the philosophy of technology. Existing work in this subfield includes research on the phenomenology of augmented experiences, the metaphysics of virtual objects, and different ethical issues associated with AR systems, including (but not limited to) issues of privacy, property rights, ownership, trust, and informed consent. This paper addresses some epistemological issues posed by AR systems. I focus on a near-future version of AR technology called the Real-World Web, which promises to radically transform the nature of our relationship to digital information by mixing the virtual with the physical. I argue that the Real-World Web (RWW) threatens to exacerbate three existing epistemic problems in the digital age: the problem of digital distraction, the problem of digital deception, and the problem of digital divergence. The RWW is poised to present new versions of these problems in the form of what I call the augmented attention economy, augmented skepticism, and the problem of other augmented minds. The paper draws on a range of empirical research on AR and offers a phenomenological analysis of virtual objects as perceptual affordances to help ground and guide the speculative nature of the discussion. It also considers a few policy-based and designed-based proposals to mitigate the epistemic threats posed by AR technology.

Keywords Augmented reality · Philosophy of technology · Augmented epistemology · Attention economy · Affordance perception · Real-World Web

This article is part of the Topical Collection on*Information in Interactions between Humans and Machines*

✉ Cody Turner
cody.turner@uconn.edu

1 Philosophy Department, University of Connecticut, Storrs, CT, USA

Augmented reality (AR) technologies function to 'augment' normal perception by superimposing virtual objects onto an agent's visual field. According to Azuma (1997), AR can be defined as any technology that is interactive in real-time, registered in 3-D, and conjoins the virtual with the physical. AR technology was first created in the late 1960s by the computer scientist Ivan Sutherland but did not enter the public consciousness until the Boeing engineer Tim Caudell coined the term 'augmented reality' in the early 1990s (Caudell & Mizell, 1992).[1] One of the most popular applications of the technology to date is the game *Pokemon Go* by *Niantic*, which has reached over a billion downloads (NintendoSoup, 2019). Pokemon Go players use a cellular map to help them find virtual Pokemon avatars pinned to various GPS locations. While Pokemon Go and other contemporary AR apps like Snapchat mandate the use of handheld smartphones, some of the most promising versions of the technology come in the form of wearable devices like *smart glasses* and even *smart contact lenses*.

The philosophy of augmented reality is a small but growing subfield within the philosophy of technology. Existing work in this subfield includes research on the phenomenology of augmented experiences (Liberati, 2018; Simonetta, 2015; Wellner, 2013), the metaphysics of virtual objects (Chalmers, 2017; McDonnell & Wildman, 2019), and different ethical issues associated with AR systems, including (but not limited to) issues of privacy, property rights, ownership, trust, and informed consent (Friedman & Kahn, 2000; Brinkman, 2014; Wassom, 2015; Wolf et al., 2015, Neely, 2019). This paper addresses some epistemological issues posed by AR systems. I focus on a near-future version of AR technology called the *Real-World Web* (RWW for short, Smart, 2012), arguing that the RWW threatens to magnify or present new versions of three existing epistemic problems in the digital age: the problem of digital distraction, the problem of digital deception, and the problem of digital divergence. While much of the discussion is necessarily speculative, it is imperative to consider ethical, practical, and epistemological problems associated with emerging technologies before these problems fully materialize in reality. This is especially true of a technology like the RWW, which promises to radically transform the nature of our relationship to digital information by mixing virtual environments with physical spaces.

A large portion of the paper is devoted to the problem of digital distraction (Hanin, 2020) which, as I explain, is primarily a manifestation of surveillance capitalism and the online attention economy (Williams, 2018). I argue that the RWW threatens to exacerbate the problem of digital distraction for two main reasons. First, the self-tracking mechanisms incorporated into the relevant AR technology will enable surveillance capitalists to collect more fine-grained personal data, which they can then leverage to more effectively capture user attention. Second, superimposed virtual objects are especially distracting in virtue of being perceptual affordances that actively 'call out' for the attention of users. I offer a phenomenological analysis

[1] Caudell and his colleague David Mizell are credited with creating the first industrial augmented reality system, which was used by Boeing factory workers to assist them in manufacturing. The system projected a digital diagram of the manufacturing process onto the external world and served as a cheaper alternative to the physical diagrams previously used to guide the Boeing factory workers. AR continues to be used in the field of manufacturing today (Nee et al., 2012).

of virtual objects in terms of J.J. Gibson's affordance-based model of perception to illustrate the latter point (Gibson, 1979).

Then, I turn to the problem of digital deception, which I understand in an inclusive sense as referring to any use of digital technologies to generate misleading or false appearances of reality. Two contemporary forms of digital deception are fake news articles and deepfake videos. I argue that RWW users may also have to deal with fake physical objects, which is to say, 3-D virtual objects that are phenomenologically indistinguishable from real physical objects. The inability to differentiate between the virtual and the physical in AR space gives rise to what Palermos (2017) calls 'augmented skepticism.' I motivate the possibility of augmented skepticism by drawing on empirical research in the field of AR and examine this possibility through the lens of my affordance-based analysis of virtual objects.

Finally, I consider how the RWW is poised to exacerbate the problem of digital divergence, which is the epistemic problem associated with the fact that digital consumers are increasingly living in different informational universes or 'filter bubbles' (Pariser, 2011, Watson, 2015). I contend that RWW users will not just inhabit filter bubbles in cyberspace but will come to occupy 'real-world filter bubbles' in virtue of having different digital information superimposed onto their perceptual fields. I distinguish between two different kinds of real-world filter bubbles (Platform-specific bubbles and Personalization bubbles) before likening the emergence of such bubbles to a version of the problem of other minds, which I call 'the problem of other augmented minds.' The problem of other augmented minds pertains to both the relationship between augmented subjects and other augmented subjects and the relationship between augmented subjects and nonaugmented subjects. I briefly explain how this problem may give rise to an 'augmented digital divide' characterized by novel forms of epistemic injustice.

The structure of the paper is as follows. Section 1 introduces the concept of the Real-World Web, motivates the idea that the advent of the RWW is on the horizon, and outlines some ways in which the RWW may improve an agent's epistemic standing. Section 2 describes the practice of surveillance capitalism and explains how surveillance capitalism and the online attention economy have engendered a problem of digital distraction. Section 3 details how the RWW threatens to magnify the problem of digital distraction by giving rise to an 'augmented attention economy.' Sect. 4 considers a few different ways to mitigate the epistemic harms of the augmented attention economy, such as the development of digital wellness AR apps, the promotion of technology-lite environments, and the ability to block augmented advertisements and possibly even augmented content from other users. Sections 5 and 6 discuss the problem of augmented skepticism and the problem of other augmented minds, respectively. Section 7 concludes.

1 The Real-World Web

Augmented reality exists near the middle of what Milgram and Kishino (1994) call the reality-virtuality continuum. On one end of the continuum is *physical reality* or reality as it is normally perceived through the senses without technological augmentation. On the other end of the spectrum is *virtual reality*, which involves total virtual immersion, or the experience of entirely exiting the physical world and entering a 3-D modeled virtual world. At the center of the spectrum lies mixed reality technologies, which encompasses both *augmented reality* and *augmented virtuality*. Augmented virtuality is essentially the inverse of augmented reality, as it concerns integrating real-world information into a virtual environment as opposed to digital information into a physical environment.[2]

Augmented reality does not yet have a ubiquitous presence in society, but the technology is poised to have a positive impact on a variety of industries beyond just entertainment, including (but not limited to) healthcare, education, retail, manufacturing, and tourism.[3] For example, in the retail space, the furniture company IKEA has created an AR app which enables consumers to project 3-D virtual models of IKEA products into their homes to assist them in making prudent purchasing and spatial design decisions. Overall, the global market for AR is incredibly diverse and estimated to increase at a compound annual growth rate of 43.8%, reaching a value of roughly 340 billion dollars by 2028, according to a report by Grand View Research, Inc.[4]

One promising consumer application of the technology is the possibility that we will create something resembling what Paul Smart has called *the Real-World Web* (Smart, 2012, 2014, 2018):

"Rather than information access requiring perceptual detachment and disengagement from our immediate surroundings (something that is required even with the most portable of mobile devices), the notion of the Real-World Web seeks to make Web-based information access a standard feature of our everyday sensorimotor engagements with the world—it seeks to make the Web part of the perceptual backdrop against which our everyday thoughts and actions take shape" (Smart, 2012: 458).

The Real-World Web represents a novel form of ubiquitous computing that could, in principle, be realized by AR technology.[5] The basic idea involves a sophisticated pair of AR smart glasses or smart contact lenses that allow for the superimposition of digital information onto the physical world in the form of three-dimensional virtual objects. I am unsure precisely how narrow Smart's conception of the RWW is meant to be, but I will choose to understand the RWW in a general sense as referring to any future AR head-mounted display or contact lens that effectively replaces the smartphone and produces a seamless and possibly even photorealistic experience

[2] Of course, there are no neat distinctions between these subclasses of technologies. As Giovanni Simonetta points out, "the more we venture towards the center of the continuum, the more difficult it is to say whether we are in AV or in AR" (Simonetta 2015: 93).

[3] In the healthcare industry, surgeons can benefit from AR by using it to project visualizations of relevant anatomy and biomedical data onto the surgical environment (Dickey et al., 2016; Feng et al., 2014). One innovative AR application in this field is Accuvein, an AR device that allows clinicians to perform real-time vein visualizations of patients.

[4] View the report here: https://www.grandviewresearch.com/.

[5] Mark Weiser (1991) refers to ubiquitous computing as a kind of *embodied virtuality*. While the contemporary internet of things (IoT) fits this description, the Real-World Web is arguably the perfect illustration of the idea of embodied virtuality.

of virtual content in physical space. The RWW, as I conceptualize it, may involve any number of reality augmentations, ranging from virtual display screens (e.g. texts threads, websites) to virtual jewelry, virtual architecture, and even real-time 3-D virtual representations of other people (Billinghurst & Kato, 2002).[6] In essence, the RWW is the fruition of Ivan Sutherland's original vision of AR as 'the Ultimate Display' which seamlessly integrates the digital world into the physical world, thereby replacing the need for browser interfaces and screen-based information displays (Sutherland, 1965).

Modern computing devices typically involve tactile inputs and visual screen-based outputs, meaning that agents request information from the Web by typing onto a keyboard and receive the information in the form of two-dimensional text or images that appear on a display screen. The RWW, by contrast, will be controlled by voice command through the use of speech recognition technology or possibly even by eye movement via eye-tracking mechanisms. Further, RWW devices will likely come equipped with motion sensors that deliver context-relevant information to users in real-time and hand gesture recognition systems that enable users to manipulate virtual objects and virtual displays in the physical world (Sharp et al., 2015). It is worth clarifying that the RWW is not synonymous with 'the Metaverse' insofar as the latter is understood as pertaining exclusively to virtual reality technology. Unlike the Metaverse, the goal of the RWW is not to escape the physical world by fully immersing human agents within cyberspace but rather to enhance the physical world by overlaying virtual content onto sensory perception.

The RWW is at this point merely hypothetical, but it is easy to see how the continuation of certain technological trends could give rise to such a device. A prelude to the idea can be found in the empirical work of Matsumoto et al. (2008), who manufactured and tested a functional prototype of what they call *the embodied web*, which "aims to provide an interface using real-world embodied interaction to provide a computer-augmented reality that accesses web services" (2008: 339). Today, various AR smart glasses or head-mounted displays can be regarded as rudimentary versions of the RWW. This includes devices like Google Glass Enterprise Edition 2, Microsoft's HoloLens 2, Sony's SmartEyeglass, Meta's Space Glasses, and the Magic Leap One. The company Apple is also in the process of constructing AR glasses which sync up to the iPhone and give users the ability to instantaneously project any information from the device (e.g. websites, videos, traffic alerts, weather updates) directly onto perception. Moreover, companies and university research teams are working on producing smart contact lenses via micro-fabrication techniques (Lingley et al., 2011). The company Mojo Vision, for example, is in the process of developing a multi-functional smart contact lens prototype called the Mojo Lens. Finally, Niantic has built a software platform for augmented reality apps called 'Lightship', which CEO John Hanke hopes will be instrumental in creating the RWW or what he calls 'the real world metaverse.' The company is currently

[6] AR technology could replace contemporary video conferencing by allowing physically remote people to feel present to the AR user. Microsoft's Holoportation has already successfully tested an AR communication device of this sort, which they call 'virtual 3D teleportation' (Orts-Escolano et al., 2016).

refining Lightship by constructing a highly detailed 3-D map of the world using a ‘Virtual Positioning System’ (as opposed to Global Positioning System).[7]

The following graphic depicts the evolution of computing technologies and illustrates how the RWW might be construed as the natural successor of contemporary smartphones: [8]

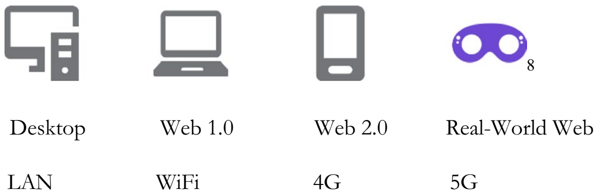

The RWW may not even be the last iteration of this evolutionary process. If smart devices continue to become smaller, more technologically advanced, and tightly coupled with cognition, we may someday develop what the philosopher Michael Lynch calls *neuromedia* (2014), which is an invasive brain-computer interface device that bypasses perception altogether and essentially implants the internet in the head by giving agents the ability to access and upload Web-based information via mental commands or the power of thought alone.

Before discussing some epistemological concerns raised by the RWW, a few qualifications are in order. First, it is vital to acknowledge that the future of technology is contingent on a plethora of societal factors and, for this reason, is very difficult to predict. In 2014 there was substantial hype surrounding the release of the Google Glass explorer smart glasses, but the device never achieved widescale uptake due to a bevy of privacy invasion concerns related to the possibility of surreptitiously recording others without their permission (Denning et al., 2014). This example demonstrates that when it comes to emerging technologies like the RWW, it is often impossible to anticipate exactly *when* and *how* the technology will be successfully implemented. The goal here is, therefore, not to engage in technological prophecy, or worse, technological fatalism. Rather, this paper understands the RWW to be a plausible technological projection based upon current trends in the development of augmented reality technology and smart devices more generally.

Of course, the actual properties of such a device are speculative because it is uncertain how the development of the technology will play out. For example, some contemporary AR devices like the Microsoft HoloLens utilize holographic projection technology (i.e. the projection of holographic images onto a physical space), whereas others like Magic Leap deploy retinal projection (i.e. the projection of light directly onto the user’s retina). Will the RWW incorporate holographic projection,

[7] For more on Hanke’s vision and why he thinks augmented reality is preferable to virtual reality, see https://nianticlabs.com/blog/real-world-metaverse/.

[8] This graphic is the intellectual property of the company Magic Leap. The original rendition can be found at the following link: https://www.magicleap.com/en-us/about. I have received express permission from Magic Leap’s CMO and Legal Team to use the graphic in this paper.

retinal projection, or some alternative display technology? What about haptic technology to simulate the experience of touch for virtual objects? Will some tech company monopolize the AR market, leading to a single sphere of augmentation that everyone plugs into? Or will there be a variety of popular AR platforms or apps competing for prominence? (Neely, 2019). These are just a few examples of important practical and technical questions that one must be privy to when thinking about the philosophical implications of future AR technology. As will be seen, many empirical considerations of this nature are brought to bear in my discussion of the RWW.

Finally, while this paper focuses on epistemological problems posed or made worse by the RWW, it does not discount the fact that there are various ways in which AR may improve an agent's epistemic standing. Smart (2012) is optimistic about the ability of the RWW to be deployed as a productive epistemic tool. Emphasizing the 'proactive' and 'perceptually immediate' nature of the technology, he says, "Instead of the retrieval of relevant information being entirely the responsibility of the human agent, the notion of the Real-World Web advocates a more intelligent and proactive Web: a Web that is capable of anticipating users' information requirements and making that information available in ways that support cognitive activity" (Smart, 2012: 458). Smart is particularly interested in the prospect that the RWW will facilitate what he calls the *Web-extended mind*, which is one application of the more general extended mind thesis (Clark & Chalmers, 1998), or the thesis that cognitive processes can 'seep out' beyond the skull to encompass technological artifacts. Smart argues that while contemporary Web-based technologies fail to satisfy the standard 'trust and glue' conditions for cognitive extension, the RWW meets these conditions in virtue of being so tightly coupled with cognitive processes. The notion of the Web-extended mind is closely related to the concept of *internet-extended knowledge* (Smart & Clowes, 2021), which in this context denotes the idea that the RWW vastly extends our knowledge base by automatically embedding Web-based information in the physical environment to help guide thought and action. This techno-optimistic vision of the RWW put forward by Smart dovetails with Jaana Parviainen's suggestion that AR technology can free us from a state of negative knowledge (or 'non-knowing') and "reveal to us the 'hidden' knowledge and information that is an inherent part of our physical infrastructure" (Parviainen, 2017: 201).

The prospect of extending an agent's knowledge capacities via AR technology brings about exciting possibilities from both cognitive therapy and cognitive enhancement perspectives. Regarding cognitive therapy, Smart (2018) and Vold & Hernandez-Orallo (forthcoming) discuss the idea of deploying AR as a memory extender to treat Alzheimer's patients who do not have reliable access to their biological memory. A RWW device could help Alzheimer's patients recognize familiar objects and people by automatically labeling them with informational displays and support their navigational competency through the superimposition of directional indicators. As a cognitive enhancement device, the RWW can promote the realization of new kinds of extended knowledge that transcend the epistemic capacities of nonaugmented individuals. For instance, AR can extend the imagination of agents by enabling them to tangibly visualize future possibilities, past occurrences, and fictional scenarios in a realistic, fine-grained manner. Architects might harness this technological capacity to create 3-D virtual models of future construction projects, whereas educators might systematically extend students' imaginations by generating convincing virtual representations of historic people and places.

As with any technology, the RWW is not intrinsically epistemically good or bad; rather, everything ultimately depends on how it is used. Bearing this truth in mind is necessary in order to avoid the follies of Techno-Utopianism and Neo-Luddism. At the beginning of the digital age many subscribed to a utopian vision of the World-Wide Web which anticipated that the technology would yield entirely positive epistemic consequences by democratizing information and fostering the widespread dissemination of knowledge. While the Web has certainly produced these epistemic goods, it has also generated a bevy of epistemic harms, such as those related to the problems of digital distraction, digital deception, and digital divergence. If the nature of contemporary Web-user interaction is any indicator of the nature of Real-World Web-user interaction, then it is likely that RWW users will also be subject to these epistemic harms. In fact, as I will argue, the RWW threatens to magnify or present new versions of digital distraction, digital deception, and digital divergence by bringing these virtual problems into the physical world. Before explaining how and why this is the case, it is necessary to briefly describe the underlying mechanisms responsible to one degree or another for all three of these problems: surveillance capitalism and the attention economy.

2 Surveillance Capitalism and the Attention Economy

Surveillance capitalism describes a relatively new type of capitalism wherein businesses, and in particular, Big Tech companies like Google, Facebook, and Amazon, collect personal data about consumers via surveillance mechanisms, use this personal data to generate prediction models concerning what products consumers might want, and then sell these prediction models to willing advertisers to make a profit (Zuboff, 2019). The practice of surveillance capitalism was pioneered by Google in 2001 when the company decided to transform the surplus behavioral data left behind by its consumers' digital footprints into personality profiles which could be sold to advertisers as a means of commercial exchange. The advertising business model that undergirds surveillance capitalism provides companies with a strong financial incentive to monopolize user attention, regardless of the negative consequences of this hijacking of attention on individuals and society. Increased user engagement brings more advertising opportunities and gives companies the leverage to charge advertisers higher prices, which means that surveillance capitalists are highly motivated to keep users plugged into their platforms for as long as possible (Lanier, 2018). This incentive structure has led to the creation of *the online attention economy*: a digital ecosystem in which consumers are bombarded by an abundance of information competing for their limited attention.

The online attention economy is especially persuasive and insidious due to its use of adaptive algorithms via machine learning mechanisms that target consumers on a granular, individual level. By monitoring users' online activity in real-time, companies can produce increasingly personalized content and advertisements that are unparalleled in their ability to capture user attention. These personalization mechanisms permeate most internet platforms, from Google's search engine to YouTube's video recommender system to Facebook's newsfeed. Furthermore, as Nir Eyal details in his book *Hooked: How to Build Habit-Forming Products*, digital design structures like the 'pull-to-refresh' button and the 'infinite scrolling' function operate in much the same way that slot machines do; namely, they offer intermittent variable rewards (in the form of notifications, likes, tweets, pictures, etc.) as a way to motivate users to remain tuned in to the platform (Eyal, 2014). Many apps even implement bright color-coding schemes designed to attract the human eye (e.g. bright red notification bubbles). Finally, even when agents are not using their smartphones, they are continually being barraged with push notifications from apps that vibrate and light up the screens of their devices, causing their attention to be diverted away from the task at hand.

The online attention economy has engendered a significant *problem of digital distraction*, wherein we are faced with "an asymmetric matchup between our fallible mechanisms of self-restraint and armies of engineers, programmers, designers, and executives working to extract ever-smaller 'slivers' of our focus" (Hanin, 2020: 3). There are at least two reasons why the problem of digital distraction is epistemic. First, increased cognitive distraction correlates with reduced cognitive efficiency, where 'cognitive efficiency' is understood as the capacity to achieve intellectual goals at an optimal rate using minimal wasted mental resources. Empirical evidence suggests that distracted agents are, other things being equal, less cognitively efficient than non-distracted agents because they bear higher cognitive costs associated with continuous task-shifting (Ophir et al. 2009). Second, and more perniciously, digital distraction can undermine a user's cognitive agency, at least when it takes the form of internet addiction (Castro & Pham, 2020). Weakened cognitive agency is, in this context, associated with not just heightened attentional capture, but a loss of attentional control. As James Williams explains in his book *Stand out of our Light: Freedom and Resistance in the Attention Economy* (2018), "the main risk information abundance poses is not that one's attention will be occupied or used up by information, as though it were some finite, quantifiable resource, but rather that one will lose control over one's attentional processes" (Williams, 2018: 37). This threat to attentional control is made salient by the fact that the online attention economy functions as a negative feedback loop: the more tech companies keep one's attention on their platforms the more personal data they can accrue, and the more personal data they accrue the more easily they can monopolize one's attention.

3 Affordance Perception and the Augmented Attention Economy

The advent of the RWW opens the door to the unsettling possibility that the mechanisms of surveillance capitalism will play an even more distracting role in the mental lives of digital consumers. In a future with the RWW, the attention economy will no longer be confined to the screens of smartphones but will seep out into the real world and invade visual perception itself. Instead of receiving vibrating push notifications that appear on the screen of one's handheld device, RWW users may receive impromptu digital advertisements and reminders that appear directly in their visual fields. Call this *the augmented attention economy*. The concern that some form of ubiquitous computing (like the RWW) will overwhelm our attentional capacities is not novel. In the early 1990s, Mark Weiser introduced the notion of *calm computing* (Weiser, 1991) in response to the worry that ubiquitous computing technologies will function to overload agents with digital information or overwhelm their perceptual processing capabilities. Weiser and Brown define a calm computing device as "that which informs but doesn't demand our focus or attention" (Weiser & Brown, 1996). The basic proposal is to design computing technologies that operate quietly in the background of perception by targeting the periphery of a user's attentional field. The central worry expressed in this section is that the RWW will serve as the antithesis of a calm computing device in virtue of incessantly demanding the focal attention of users.

There are at least two different ways in which the RWW is poised to exacerbate the problem of digital distraction. First, *the self-tracking mechanisms incorporated into the relevant AR technology will enable surveillance capitalists to collect greater (and more fine-grained) amounts of personal data that they can then leverage to capture user attention more effectively*. Surveillance capitalism has already expanded beyond the online realm and into the real world with the advent of *the internet of things (IoT)*. It is not just 'browsing' data (i.e. data extracted from the digital footprint agents leave behind on the internet) that is being tracked, analyzed and sold by tech companies. Many agents are now equipped with or surrounded by an array of smart devices (e.g. smartphones, smart watches, smart houses) that actively track their behavior in the physical environment and collect 'real-world' personal data (e.g. biomedical data) in the service of surveillance capitalism. Advanced AR technologies contain various vision-based and motion-based sensors that surpass the self-tracking capabilities of contemporary smartphones. In order to generate three-dimensional virtual objects that appear to be embedded in physical space, the user's head position must be actively tracked by the AR device. This is typically accomplished via a combination of hardware and software components like depth cameras, gyroscopes, accelerometers, global positioning systems (GPS), eye-tracking mechanisms, and simultaneous localization and mapping (SLAM) software (Bostanci et al., 2013; Liu et al., 2016). There is even ongoing research into the development of 'AR cloud-based tracking' according to which "features captured by a user's device are uploaded to the cloud and fused to provide a ubiquitous tracking service" (Billinghurst, 2021: 2). AR cloud-based tracking could conceivably lead to the existence of AR smart cities characterized by pervasive surveillance mechanisms. The increasing integration of such self-tracking technologies into our daily lives threatens to magnify data privacy concerns and empower tech companies to engage in even more persuasive forms of targeted advertising. For instance, eye-tracking data could be used by advertisers to determine engagement down to where a user is looking, and physiological data could be employed to gauge a user's emotional state, so advertisers know when the user is in the most persuadable frame of mind. All of this is compounded by the fact that the advent of the RWW enables the possibility of personalized advertising in physical space and not just cyberspace. In the future, RWW users might walk into a store and immediately be bombarded with personalized virtual recommendations that appear in their visual fields and direct them to different items in the store that most align with their personality profiles and preestablished preferences.[9]

The second way the RWW threatens to magnify the problem of digital distraction has to do with the fact that *virtual objects are themselves especially intrusive and distracting from a phenomenological perspective.* I submit that virtual objects are kinds of perceptual affordances that phenomenologically 'call out' for the attention of agents. In what follows, I introduce the concept of an affordance, analyze the phenomenology of virtual objects in terms of affordance perception, and illustrate how this affordance-based analysis sheds light on how the RWW is disposed to magnify the problem of digital distraction.

Affordances are dispositional properties that offer different courses of action to a perceiving subject. For example, a mug has the property of being graspable and an apple has the property of being eatable. In both of these instances, objects in the environment afford certain opportunities for action when presented in visual perception: the mug 'calls out' to be grasped, and the apple 'calls out' to be eaten. The notion of affordance derives from the ecological psychologist J.J. Gibson (1966, 1979), who argues that we are perceptually sensitive to affordances, and that perception itself can be explained in terms of affordance detection. In other words, Gibson asserts that affordances are both phenomenologically and metaphysically pertinent to perceptual processing. This position is reflected in Gibson's ecological theory of perception, which is a theory of *direct* perception according to which perceptual content is located not inside the skull but rather in environmentally based affordances.[10]

The affordance model of perception naturally lends itself to a phenomenological analysis of AR systems because virtual objects projected onto visual perception can afford opportunities for action in much the same way physical objects do. For example, imagine that a RWW user superimposes directional indicators onto their visual field to help guide them to a desired location. These directional indicators can be construed as perceptual affordances because they *afford* the opportunity to walk in the correct direction to the destination. One might object that the 3-D indicators are not themselves affordances but instead *serve as a means* to help the user see other physical objects in their environment as affordances. For instance, the directional indicators might enable the user to perceive a sidewalk that they previously did not attend as 'walkable'. While this analysis may hold for 3-D directional indicators and other virtual tags that provide information about physical objects and locations, there are nevertheless instances in which virtual objects are clearly perceived as affordances in their own right. Imagine that a RWW user projects onto their

9 Sandor et al., 2015 propose light field displays as a possible means for creating a realistic AR experience in the absence of self-tracking mechanisms.

10 According to Gibson, physical stimuli are imbued with informational content prior to any internal representation, and perception is primarily a matter of directly detecting this informational content in the stimuli. The salient information detected on the view is information concerning what courses of action external objects afford (i.e. the theory says that we primarily perceive objects through their affordances).

visual field virtual representations of common household items. In one instance, they superimpose a life-like virtual chair onto their perception of the world, and in another a virtual cup. It is reasonable to assume that the user will perceive the chair as sit-able and the cup as grasp-able even though these items are not actual physical objects. Gibson himself recognizes that humans have constantly and strategically added new affordances to the environment and manipulated the external world so as to make "more available what benefits him and less pressing what injures him" (Gibson, 1979, 130). This is no less true in the case of virtual objects being superimposed onto the physical world by an AR device.

I am not the first scholar to apply the concept of affordances to virtual objects. Affordances have played a significant role in virtual reality research for decades (Flach & Holden, 1998). Regia-Corte T. et al. (2012), for example, demonstrate that it is possible to quantify the perception of affordances in virtual environments, whereas Grabarczyk and Pokropski (2016) argue that perceiving virtual objects as affordances is necessary for producing the experience of presence and immersion in such environments. Simonetta (2015) is, to my knowledge, the only other scholar to explicitly apply the affordance model of perception to augmented reality systems (as opposed to virtual reality systems).[11]

In order to understand how this phenomenological analysis of AR sheds light on how the RWW is disposed to exacerbate the problem of digital distraction, two further clarifications must be made concerning the concept of affordances. First, there are two different kinds of affordances: *physical* affordances and *mental* affordances. Thus far, I have been referring solely to physical affordances or perceived opportunities for physical action such as walking, eating, and grasping. However, as Tom McClelland (forthcoming) explains, it is also possible to speak of mental affordances, which can be defined as perceived opportunities for intellectual action (e.g. a scientific hypothesis affords doubting, a familiar place affords recognition, a captivating book affords reflection). Second, affordances do not just *present* opportunities for action but also *motivate* or *solicit* specific courses of action. When I have a perceptual experience of a chair as sit-able, I do not merely experience the action of sitting as being open to me but also experience the chair as soliciting or inviting that particular course of action. Somehow the perception of the affordance compels me to initiate the action that is afforded.[12] McClelland (2019) explains soliciting affordances in terms of what is called 'affordance potentiation', the idea that "when we perceive an affordance, the motor patterns responsible for performing the afforded action (or parts of that action) are automatically readied" (McClelland, 2019: 8). McClelland draws on a range of empirical studies which support the notion of

[11] In his paper 'The Realism and Ecology of Augmented Reality', Simonetta contends that AR objects are best thought of as environmental affordances, and then uses this insight to argue for virtual realism, the idea that virtual objects are metaphysically real entities: "The main argument was that these [AR digital objects] are indistinguishable from real objects insofar as their affordances are exactly the same as those of real objects, i.e., they motivate the very same users' behaviors that would have been motivated by usual perceptions" (Simonetta 2015: 108).

[12] The concept of 'soliciting affordances' is commonplace in the literature and is widely acknowledged as a standard feature of affordance perception (Siegel 2014).

affordance potentiation and proposes that this motor phenomenon may underpin the phenomenology of soliciting affordances.

Virtual objects can phenomenologically function as both physical affordances and mental affordances, meaning that they can afford physical actions and intellectual actions. The foremost intellectual action afforded by virtual objects is arguably the intellectual act of attention.[13] Given that affordances possess motivational force, this means that virtual objects superimposed onto an agent's visual field by an AR device literally solicit one's attention. In other words, virtual objects do not *passively* wait to be perceived by the RWW user but instead *actively* 'call out' for the user's attention in virtue of being perceptual affordances. The relevance of this phenomenological analysis to the problem of digital distraction should now be coming into view. If RWW users are, on average, as absorbed with and dependent on different elements of the infosphere as contemporary internet users, then they will likely tend to have visual fields that are tiled over with an array of virtual objects. The worry is that a visual phenomenology characterized by a constant influx of 'digital affordances' will exacerbate the problem of digital distraction by actively pulling one's attention in directions that are inconducive to their intellectual goals.

To see the concern more clearly, it may be helpful to compare the phenomenological predicament of a RWW user with the phenomenological predicament of a contemporary smartphone user. Smartphones can also be conceptualized as perceptual affordances, as they are typically perceived as 'usable' in the sense that they solicit the physical action of picking up and using the device. This affordance-based analysis of smartphones might partially explain recent empirical findings which suggest that simply having a smartphone in one's immediate vicinity impairs cognitive efficiency and the ability to focus on tasks (Ward et al., 2017). It would make sense that the mere presence of a smartphone in one's visual field undermines cognitive efficiency given that smartphones phenomenologically function as objects of temptation, meaning that, as perceptual affordances, they 'call out' to be used. It is also reasonable to think that, due to the prominent role smartphones play in our lives, the affordances offered by them have an especially high salience compared to other physical affordances in the environment.

The RWW user will presumably not have a physical phone in close proximity that 'calls out' to be used, but their perceptual field may be polluted with a bevy of virtual objects that 'call out' be attended to or that solicit the intellectual action of attention.[14] Simply put, the RWW changes the motivational landscape of perception in a way that the mere presence of a handheld smartphone does not. While smartphones may be distracting in virtue of being perceptual affordances, their presence in perception is nothing compared to having a visual field littered with 3-D virtual

[13] Of course, virtual objects can afford other intellectual actions as well. Consider a pair of AR glasses that projects aesthetically pleasing digital artworks onto an agent's visual field as the agent walks around a city. Such virtual objects may afford a variety of intellectual actions, including imagining, reflecting, and evaluating. However, this does not negate the fact that the foremost intellectual action afforded by virtual objects is the act of attention.

[14] The following video, entitled 'Hyper Reality', illustrates what the phenomenology of the RWW might look like: https://vimeo.com/166807261.

objects. The ability to project virtual objects onto a user's visual field means that AR smart glasses (or contact lenses) can more directly engender states of attentional capture than handheld smartphones. All else being equal, then, RWW users will be more vulnerable to cognitive distraction than contemporary digital consumers because of their distinctive visual phenomenology. In the worst-case scenario, one can imagine the augmented attention economy facilitating states of perceptual overload which further undermine digital consumers' capacity for attentional control and perhaps even cause them to exhibit symptoms resembling attention-deficit hyperactivity disorder (ADHD).

Notably, the problem of digital distraction is not merely epistemic but also poses real-world safety risks (Pase, 2012). It is not uncommon to see contemporary smartphone users bump into physical objects in public because they are too distracted by the device and thus oblivious of their surroundings. During the Pokemon Go craze in the summer of 2016, for example, there were news stories about players being injured because they were too absorbed in the game and not attentive enough to the external environment. In the case of the RWW, the superimposed virtual objects may not just *distract agents from* the physical world, but also *obstruct their perception of* the physical world. Digital information which is overlaid onto an agent's visual field might in some circumstances function as a kind of *virtual veil of perception* that prevents the agent from gaining full experiential access to their physical surroundings (e.g. a virtual display screen blocking one's view of nature). This possibility of digital obstruction, and not only of digital distraction, renders use of the RWW considerably more conducive to bodily injury than use of contemporary smartphones.

4 Digital Wellness and Augmented Ad Blocking

A skeptical reader might think that the worries raised in the previous section surrounding the RWW and the augmented attention economy are speculative at best and exaggerated or baseless at worst. There are at least two key questions that the skeptic might raise in this context. First, even if the RWW becomes widespread, why think that agents will have visual fields that are 'polluted' with virtual objects? Perhaps agents will use the RWW sparingly or at least be judicious enough not to overload their perception with virtual content. Second, even if RWW users overload their perception with virtual content, why think that this will necessarily endanger their capacity for attentional control? Many people in the contemporary world live in highly distracting city environments but learn how not to be overwhelmed by such environments through exposure and habituation. As RWW users become accustomed to the technology, and eventually become *augmented reality natives* (i.e. people who are born into a world in which AR technology is widespread), they may, in a similar manner, learn how to navigate augmented spaces without becoming unduly distracted.

The skeptic is, of course, correct that the worries surrounding the augmented attention economy are at this point speculative, but these worries are certainly not baseless. Quite to the contrary, they are based upon current trends of surveillance capitalism and digital distraction. They are based upon the assumption that the machine learning methods, persuasion tactics, and financial incentive structure underlying the modern attention economy will continue to drive the augmented attention economy. The relevant claim is simply that the problem of digital distraction is, *all else being equal*, poised to become worse in the age of the RWW than it currently is in the age of the smartphone. As discussed in the previous section, this exacerbation concern stems largely from the development of more precise self-tracking techniques and the phenomenological analysis of virtual objects as perceptual affordances which solicit the intellectual action of attention (i.e. 'call out' for one's attention in perception). Given this phenomenological analysis, and the continuation of current trends regarding surveillance capitalism, the prospect that RWW users will manage to avoid significant digital distraction seems relatively low.

Assuming that there is reason to be concerned that the RWW will magnify the problem of digital distraction, it is worth thinking about what measures can and should be taken to address the problem. One possibility is that tech companies will self-regulate in response to public pressure campaigns. For example, the former Google design ethicist Tristan Harris has started a movement called 'Time Well Spent' which pressures tech companies to consider consumers' well-being and values when constructing their algorithms, design layouts and business models (Harris, 2018). Instead of aiming to hijack consumer attention at all costs, Harris urges tech companies to construct their platforms such that consumers are incentivized to use them in a way that they would retroactively consider 'time well spent.' One promising recent development that aligns with Harris' vision is the emergence of so-called *digital wellness technologies*, which "use the same attractive qualities of other persuasive apps to motivate users towards behaviors that are personally and socially valuable, such as exercise, wealth-management, and meaningful communication" (Specker Sullivan & Reiner 2021: 413). Some digital wellness technologies, such as Apple's downtime feature or the app 'Moment', focus on promoting attentional well-being by motivating users to engage in less daily screen time. The app 'Moment' does this by actively tracking a user's screen time and sending systematic alerts notifying the user when to take a break from the phone. It is easy to imagine an AR equivalent of this app that RWW users might deploy to help retain attentional control. One might, for example, decide to preprogram their RWW device to superimpose virtual alerts onto their visual field as soon as they surpass one hour of allotted social media time. These virtual alerts will be incredibly distracting, but in this context, they will function to distract from cognitive distraction itself by soliciting attention away from epistemically unproductive pursuits and towards pursuits that further one's intellectual goals.

Despite the emergence of digital wellness technologies, many are rightly skeptical that the problem of digital distraction will ever be effectively managed via corporate self-regulation. Essentially, in the absence of completely reforming the advertising business model, tech companies will never be financially incentivized to help secure the attentional well-being of digital consumers. Quite to the contrary, they will always be incentivized to capture user attention for the sake of profit maximization. Recognizing this fact, some scholars have advocated for government-based regulatory strategies, ranging from mandatory product safety labeling of smart devices (Castro & Pham, 2020) to the advancement of what Hanin (2020) calls *technology-lite environments*, which in the case of the RWW, might involve restricting the use of reality augmentations in certain public venues like schools and public parks.

Another way to empower the individual RWW user to resist the distracting forces of the augmented attention economy is to provide them with the technological ability to block unwanted augmentations from being superimposed onto their fields of perception. We might refer to such an ability as *augmented ad blocking.* Do RWW users have a right to equip themselves with augmented ad blockers as a way to promote attentional well-being? The ethics of ad blocking is notoriously controversial. Some people argue that installing ad blockers is unethical because it robs digital platforms of ad revenue and is therefore, like music piracy, analogous to theft. By contrast, Zambrano and Pickard (2018) contend that ad blocking is ethical on the grounds that it is analogous to other commonly accepted kinds of ad avoidance behavior (e.g. flipping through TV channels during commercial breaks). I contend that the rationale for the ethical legitimacy of ad blocking is considerably stronger in the case of the RWW than it is for the contemporary Web because of the especially intrusive nature of augmented advertisements. As Wolf et al. (2015) observe, AR users arguably have ownership rights over their visual experience and should thus be able to control which augmentations impinge upon their perception: "there is a reasonable argument that an individual can choose to use an AVFD [augmented visual field device] 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" (Wolf et al., 2015: 130).

In addition to augmented advertisements, RWW users may also be confronted by virtual objects projected onto their perception by other AR users. Does the right to control one's own visual experience entail that RWW users should be able to block augmentations from other AR users and not just AR advertisers? On the one hand, the ability to block unwanted virtual content from other AR users might be seen as ethically analogous with the ability to block or mute unwanted tweets from other Twitter users. However, one could argue that the 'in-person' aspect of the RWW morally complicates the question of whether it is permissible to technologically silence other AR users in this way. The concept of 'silencing' in feminist epistemology and philosophy of language is regarded as a kind of *testimonial injustice* wherein a speaker is harmed in their capacity as a knower in virtue of not being able to fully or accurately express their thoughts in verbal discourse, typically for reasons having to do with power hierarchies and group identity (Fricker, 2007).[15] Is there a case to be made that technological silencing may, in some circumstances at least, constitute a type of testimonial injustice? Specifically, could a RWW user be considered a victim of testimonial injustice if they are prevented from effectively expressing themselves by other RWW users who block or silence their reality augmentations? The answer to this question surely depends on various factors, including circumstantial details about relevant power differentials and whether the

[15] There are three different kinds of silencing: locutionary silencing, illocutionary silencing, and perlocutionary silencing. An interesting question is which of these three kinds of silencing (if any) 'technological silencing' falls under.

superimposition of virtual content can broadly be construed as a type of speech act.[16] While it is beyond the scope of this paper to adjudicate these matters, doing so may be necessary to determine the ethical permissibility of blocking virtual content from other AR users.[17]

The past two sections have investigated some epistemic and practical risks associated with reality augmentations that distract us from or obstruct our view of the physical world. An opposite epistemic problem arises when virtual objects seamlessly blend into the external environment so as to be indistinguishable from physical objects. It is this problem to which I now turn.

5 Augmented Skepticism

In addition to magnifying the problem of digital distraction, the RWW is also poised to engender a new type of *digital deception*. As stated in the introduction, I understand 'digital deception' in a broad sense to include any use of digital technologies to generate false or misleading appearances of reality. Three contemporary forms of digital deception are fake news articles, fake images (i.e. doctored photographs), and deepfake videos. 'Fake news' refers to false or misleading news stories that are presented as real, whereas 'deepfakes' involve the use of machine learning technology to generate synthetic videos or images that appear real (Fallis forthcoming). The prevalence of these forms of digital deception can be at least partly (if not largely) explained in terms of the mechanics of the online attention economy. Whether it is a Photoshopped image on Instagram, a deepfake video of a celebrity, or a fake news story that provokes political outrage, digital deception is an effective means of capturing user attention. This is reflected by the fact that recommendation algorithms tend to prioritize sensational content because such content optimizes user engagement. AR technology threatens to bring the problem of digital deception into the physical world, as it is capable of producing virtual objects that are phenomenologically indistinguishable from real physical objects. Empirical researchers have already developed a variety of techniques to create life-like reality augmentations, including shadow representation of virtual objects (Sugano et al., 2003), real-time ray tracing (Santos et al., 2012), image-based lighting and environment illumination maps (Agusanto et al., 2003), and physical object occlusion (Breen et al., 1996).[18]

[16] Neely (2019) considers the question of whether reality augmentations should be conceptualized as speech acts in the context of discussing the ethics of augmenting public property: "If we have a national park or a public square, what kinds of augmentations would be ethical? Who should be permitted to augment the space? I see two possible answers, depending on whether we see augmentation more like speech or more like graffiti" (Neely, 2019: 17)

[17] What if a RWW user had the technological capacity not only to block virtual content from other users, but literally silence other users by muting them such that one can no longer hear them speak? This kind of literal technological silencing would assuredly constitute testimonial injustice in certain circumstances.

[18] Sandor et al., 2015 claim that 'true augmented reality' necessarily involves realistic virtual objects which seamlessly blend into the background physical environment. They propose an *AR Turing Test*, which could be implemented in the future to determine whether a system qualifies as True AR.

These empirical developments suggest that RWW users will have to be on guard against not just fake news stories and deepfake videos, but also fake objects, which is to say, virtual objects that present themselves in perception as real physical objects with physical properties (physical location, colors, etc.).

Of course, virtual realists like Chalmers (2017) would take issue with the idea that virtual objects are 'fake.' There are two main theories concerning the metaphysical status of virtual objects: *virtual realism* and *virtual fictionalism*. Virtual fictionalism holds that virtual objects are fictional objects (see McDonnell & Wildman, 2019) whereas virtual realism claims that virtual objects are real, mind-independent entities with causal powers. If virtual fictionalism is true and virtual objects do not really exist, then the perception of such objects can be regarded as a kind of hallucination. By contrast, if virtual realism is true, then the perception of virtual objects will be a kind of illusion insofar as the objects are mistakenly perceived as physical instead of as virtual (Chalmers, 2017, forthcoming).[19]

It is true that certain augmented aspects of experience will be clearly distinct from physical reality and therefore non-illusory. Examples of virtual objects that will be easily identifiable as such include superimposed text and videos, fluorescent virtual tags hovering above physical objects, and navigation arrows. In all of these instances, the RWW user will have no trouble distinguishing what is physical from what is virtual. Other components of augmented reality, however, may be more epistemically problematic, for it is reasonable to assume that there will be significant consumer demand for life-like virtual objects. For example, users of an AR shopping app will presumably want the virtual clothes that they try on to look as realistic as possible. Moreover, some agents may want to not just try on virtual accessories to help them decide which physical accessories to buy, but actually use virtual accessories as a vehicle for self-expression and identity affirmation. Instead of physical makeup, body piercings, jewelry, and tattoos, agents may choose to wear life-like virtual counterparts of these items, which will be publicly observable to anyone equipped with the relevant AR technology. As Wolf et al. (2015) envision: "Those viewing someone who is virtually decorated through an AVFD [augmented visual field device] will see that person as part physical, part holographic and perhaps be unable to distinguish between the two" (130). Importantly, the worry is not only that agents may be *accidentally misled* by virtual objects in augmented reality space. There will also be an incentive to use AR technology to *intentionally deceive* others for financial gain or other ethically suspect reasons. For instance, Neely (2019) considers the possibility of a real estate agent augmenting a piece of property to make it look more pristine than it actually is, and then overpricing the property on this basis.[20]

[19] Hallucinations involve the perception of non-existent objects, whereas illusions involve the misperception of existent objects. This is why the hallucination view of virtual objects is supported by virtual fictionalism, whereas the illusion view of virtual objects is supported by virtual realism (see Chalmers, 2019).

[20] See section 6 for some examples of how digital deception in AR space might function to amplify political polarization or even promote genocide in the worst-case scenario. These examples are presented in the context of a discussion about digital divergence, but they also fall into the category of digital deception.

The idea that RWW users will often be unable to differentiate between physical objects and virtual objects in AR space gives rise to what Palermos (2017) calls *augmented skepticism.* Augmented skepticism is not, like traditional external world skepticism, the radical thesis that it is impossible to accrue knowledge of the physical world. Rather, it can be understood as the moderate claim that it is significantly harder to achieve knowledge of the physical world in the context of a photorealistic augmented reality environment. Technologically nonaugmented subjects can acquire justified true beliefs about everyday physical objects simply by perceiving these objects in normal lighting conditions and in the absence of any relevant defeaters. RWW users, by comparison, may be accustomed to being in the presence of photorealistic virtual objects that are indistinguishable from physical objects. Insofar as this is the case, they will have to conduct further 'background checks' in order for their everyday beliefs about the physical world to be epistemically justified. These background checks may consist in haptic feedback tests (in which a user makes contact with an object to see whether it offers physical resistance) or visual tests (in which a user walks around an object to view it from different angles). As Palermos says, "In the absence of such additional background checks, 'augmented skepticism' would ensue, making it impossible to distinguish between virtually any aspect of augmented and physical reality. Perceiving and interacting with the external world would no more be the same, bringing about a dramatic change to our everyday epistemic practices" (Palermos, 2017: 143).

The affordance-based analysis of AR discussed in Sect. 3 can help us better understand the threat of augmented skepticism; particularly, why augmented skepticism poses not just an epistemic threat to RWW users but also a physical safety risk. One key fact about affordance perception is that it is possible to *misperceive* affordances. Physical objects sometimes present opportunities for action when no such opportunities exist, like when a flimsy tree is misperceived as climbable (Gibson, 1979: 139). Virtual objects are considerably more likely than physical objects to be misperceived as affording opportunities for physical action, especially if the virtual objects bear a close resemblance to their physical counterparts. For example, a lifelike virtual couch will afford the physical action of sitting even though the couch cannot actually be sat upon. Misperceptions of this sort can lead to grave bodily injury. Moreover, the concern is not just that RWW users will treat virtual objects as physical (e.g. trying to sit on a virtual couch and falling through it), but also that they will treat physical objects as virtual (e.g. walking into a physical wall on the false assumption one could walk through it). Both of these physical hazards will exist for RWW users who are confronted with the problem of augmented skepticism.

In response to this safety concern, it is important to acknowledge the possibility that future AR systems will include haptic feedback technology that enables users to touch and pick up virtual objects and perhaps even neural stimulation technology that enables them to taste and smell such objects. There is already work being conducted on creating multisensory AR systems of this nature (see Narumi et al., 2011 for taste-based applications and Pezent et al., 2019 for touch-based applications). If RWW users can pick up virtual cups, sit in virtual chairs, and taste virtual food, they will not be as in danger of falling victim to the aforementioned safety risks. That is, in the case of a highly sophisticated multisensory AR system, there may be no misperception involved when an agent encounters a virtual cup that affords the action of grasping. This perceptual experience will still be illusory to the extent that the agent incorrectly perceives the cup as a physical object instead of as a virtual object, but there will no longer be an illusion associated specifically with the relevant affordance perception. For in this futuristic scenario, the virtual cup does afford the physical action of grasping.

One obvious way to mitigate the risk of augmented skepticism is to proactively design virtual objects that are clearly distinguishable from physical objects. Palermos suggests that virtual objects “should be delineated with fluorescent borders, have a see-through effect or both” (Palermos, 2017: 146). According to Palermos, codifying AR design provisions like these into law would take the burden of risk mitigation off the shoulders of AR users. Instead of constantly engaging in cognitively demanding background checks to ensure that they are not in the presence of deceptive reality augmentations, AR users could rest assured that any virtual objects in the environment will be clearly perceivable as such. Having explained how the RWW is poised to exacerbate the problems of digital distraction and digital deception, I now consider the final problem of digital divergence.

6 Real-World Filter Bubbles and the Problem of Other Augmented Minds

Digital divergence, as I understand it, refers to the fact that digital consumers are increasingly living in different informational universes or ‘filter bubbles’ (Pariser, 2011, Watson, 2015). The problem of digital divergence is also largely a consequence of surveillance capitalism and the incentive structure underlying the attention economy. The issue is not merely that digital consumers inhabit different online platforms (i.e. some agents spend most of their time on Facebook and Twitter whereas others congregate on Reddit and YouTube). Even when occupying the same platform, users are not living in the same informational universe because personalization algorithms largely determine the content they are being fed. As explained in Sect. 2, tech companies actively track users’ online behavior via their digital footprint to deliver hyper-individualized content and advertisements that align with their preferences.

The phenomenon of digital divergence presents a variety of epistemological concerns. On an individual level, filter bubbles can facilitate confirmation bias and perhaps even lead to the inculcation of epistemic vices like close-mindedness and intellectual arrogance (Nguyen, 2020). On a societal level, they can threaten the project of deliberative democracy by undermining civic discourse and disintegrating any sense of a shared cultural reality (Bozdag & van den Hoven 2015). As Pariser says, “Democracy requires citizens to see things from another’s point of view, but instead we’re more and more enclosed in our own bubbles” (Pariser, 2011: 8). While this is an epistemologically depressing state of affairs, we can at least take solace in the fact that digital consumers still occupy the same shared physical reality, a reality that they must return to once they put down their smartphones or log off of their computers.

By mixing physical reality with virtual reality, the RWW promises to literalize the filter bubble metaphor such that people no longer share any common world. RWW users will not just inhabit filter bubbles in cyberspace but will come to occupy 'real-world filter bubbles' in the sense that they will experience the physical world differently from one another in virtue of having disparate virtual content superimposed onto sensory perception. A distinction can be made between two kinds of filter bubbles that might exist in AR space: (1) platform-specific AR bubbles and (2) personalization AR bubbles.

Regarding (1): the RWW may be spread out across multiple popular AR platforms or apps. In his forthcoming book *Reality+: Virtual Worlds and the Problems of Philosophy*, David Chalmers imagines a future in which there are numerous AR platforms available to digital consumers, such as Apple Reality, Facebook Reality, and Google Reality. He claims that each of these AR platforms will likely contain an array of virtual objects that can only be accessed by agents plugged into the relevant platform. Virtual objects of this nature can be called *public virtual objects* because they are publicly (albeit exclusively) accessible to any agent equipped with the pertinent AR technology (e.g. virtual traffic lights, virtual graffiti, virtual monuments, virtual jewelry). Chalmers says that "In Facebook Reality, there may be a virtual piano at a certain location in Washington Square. In Apple Reality, there may be a virtual sign at the same location. In Google Reality, there may be nothing there at all" (Chalmers, forthcoming: 225). To the extent that RWW users inhabit distinct AR platforms, they will have divergent experiences of the physical world and, in particular, interact with different public virtual objects. However, the existence of platform-specific AR bubbles is contingent upon the future development of AR technology. Neely (2019) discusses the possibility of living in a single augmented sphere instead of a world in which augmentation transpires over various apps. This vision of a single augmented sphere could come to pass if some tech company monopolizes the AR market such that RWW users are all plugged into the same AR platform.

Even if RWW users occupy the same AR platform and thereby experience the same public virtual objects, they will still have divergent augmented experiences insofar as their visual fields are also populated with *private virtual objects*, or virtual objects that are solely visible to the individual RWW user. Private virtual content is not restricted to virtual display screens containing personal data (e.g. texts, social media notifications, directional indicators, daily reminders, push notifications) but may also include novel augmentations of the physical world that only the individual user experiences. A RWW user may, for example, have the capacity to alter their subjective experience of the weather or seasonal conditions by superimposing clouds onto perception or by modifying the coloration of leaves to simulate the appearance of Fall. While personalized augmented experiences like this are apt to be enriching from an aesthetic, epistemic, and practical perspective, there is also a serious concern that real-world filter bubbles will exacerbate the epistemic and political problems associated with the phenomenon of digital divergence. If the RWW is governed by the same personalization algorithms that currently fuel the practice of surveillance capitalism, then users of AR technology will increasingly inhabit custom realities wherein their ideological worldview is literally reflected in their perception on the physical world. Consider a RWW user who, for political reasons, is predisposed to believe that city X is crime ridden and highly polluted. There are various ways in which the user's AR device might algorithmically reinforce this belief. When traversing the city, the user might receive continuous virtual updates in their visual field containing information about recent crimes that have been committed at their physical location. Their AR device might also come equipped with a *reality focuser function* (e.g. the iPhone's new Cinematic mode), which in this context is deployed to accentuate the run-down aspects of the city. One can imagine all the nearby litter and graffiti being brought into full focus as the RWW user walks down the street, creating the phenomenological impression that the city really is in bad shape. Alternatively, consider a RWW user who, for opposite political reasons, is predisposed to believe that city X is relatively crime free and pollution free. Their AR device might instead activate a *reality blocker function* which systematically conceals the run-down aspects of the city. In this case, one can imagine all the nearby litter and graffiti being hidden behind virtual blinders as the RWW user walks down the street, creating the phenomenological impression that the city really is in pristine condition. This example illustrates one way in which real-world filter bubbles threaten to corroborate individual cognitive biases and amplify political polarization. Essentially, when common physical reality is replaced by customized augmented realities, the prospect for a coherent social epistemology breaks down.

I submit that the real-world filter bubble concern can be likened to a version of the problem of other minds. The problem of other minds is an epistemological problem that derives from the fact that we can never directly experience the minds of other agents but instead merely infer that they have minds based on their intelligible behavior. A distinction can be made between two basic versions of the problem, which I call 'the existence question' and 'the content question':

The Existence Question: Given that I only ever directly experience the contents of my own mind, how do I know that other agents have minds at all and are not philosophical zombies?

The Content Question: Given that I only ever directly experience the contents of my own mind, how do I know that other agents experience reality in the same way I do and do not have something like inverted qualia?

Philosophical zombies are imaginary creatures that are physically and behaviorally indistinguishable from humans, but that completely lack phenomenal consciousness or inner experience (Chalmers, 1996). To question whether other agents are philosophical zombies is, in effect, to doubt the existence of their minds. The content question, by contrast, presupposes the existence of other minds and proceeds to wonder whether the nature of other minds is radically different from our own. The concept of inverted qualia (or the inverted spectrum) refers to the thought experiment in the philosophy of mind in which two individuals systematically experience the world as having different colors despite possessing the same color vocabulary (Shoemaker, 1982). For example, if I experience a tomato as phenomenally red and you experience a tomato as phenomenally blue even though we both use the term 'red' to describe the tomato, then you have inverted qualia. To question whether other agents have inverted qualia or anything like inverted qualia is to pose the content question.[21]

The real-world filter bubble concern poses the following variation of the content question, which I call *the problem of other augmented minds:*

The Problem of Other Augmented Minds: Given the ubiquity of augmented reality technology and the Real-World Web, how do I deal with the fact that other agents experience reality differently than I do in virtue of having distinct augmented qualia?

By 'augmented qualia', I just mean perceptual experiences that are technologically enhanced by AR devices. Augmented qualia could, in principle, be inverted qualia, for AR technology is capable of changing the coloration of one's visual experience. The epistemological worry surrounding augmented qualia, however, is much broader than the possibility that other agents have different color experiences than we do. Given the myriad of possible augmentations on offer, there is no telling how different an AR user's experience of the world might be from our own. Unlike the traditional problem of other minds, the problem of other augmented minds is not a harmless philosophical thought experiment but a serious practical concern with unforeseen consequences. The issue is not merely that *it is possible* that other agents have divergent experiences of reality because we only ever directly experience their overt behavior, but rather that *it will actually be the case* that other agents have divergent experiences of reality given the ubiquity of augmented reality technology and real-world filter bubbles. In addition to likely amplifying political polarization and reinforcing cognitive biases, this predicament will, at the very least, lead to inconveniences in everyday life, as agents with distinct augmented qualia may exhibit behavior that seems strange or even alien to us since they will be interacting with virtual content that eludes our perception. The *Black Mirror* episode *Men Against Fire* demonstrates how one could intentionally game the problem of other augmented minds towards malicious ends. The episode follows a group of soldiers implanted with an augmented reality device called MASS which, unbeknownst to them, alters their perception of a particular ethnic group of people. These people appear to the soldiers as humanoid roaches. The soldiers are instructed to kill the roaches, unaware that they are actually committing genocide against innocent human beings. I recognize that the scenario depicted in this episode involves a unique form of technological manipulation, and I am not suggesting that the RWW will likely give rise to dystopian outcomes of this nature. The episode *Men Against Fire* simply serves as an extreme example of the possible dangers associated with a technologically augmented world in which agents inhabit different realities.

Furthermore, the episode demonstrates that the problem of other augmented minds does not just pertain to the relationship between augmented subjects and other augmented subjects but also concerns the relationship between augmented subjects and nonaugmented subjects. This is an important point because even if the RWW becomes widespread, there will still probably be a significant number of people who

[21] Philosophical zombie arguments are deployed in the philosophy of mind to argue against physicalism about consciousness. Inverted qualia thought experiments, by contrast, are used in the philosophy of mind to argue against functionalism about consciousness.

are not equipped with the pertinent AR technology. While some agents may choose not to augment their minds for personal ideological reasons (e.g. Neo-Luddism), others may be unable to do so for socioeconomic or political reasons. This raises the possibility that the advent of the RWW will exacerbate *the digital divide*, which is to say, the gap between people who reap the technological benefits of the digital age and those who do not (van Dijk, 2006). Some agents are too poor to afford a reliable smartphone, whereas others live in countries with heavy censorship that prevent unfiltered internet access. The RWW poses the troublesome possibility that the digital divide will come to be represented by disparate perceptual experiences of reality between the haves and the nonaugmented have-nots. This 'augmented digital divide' could create or at least amplify pressing issues of epistemic and social injustice, especially if the RWW becomes indispensable to participation in society, much like the smartphone is now. For instance, it is conceivable that the RWW gives rise to a distinctive form of hermeneutical injustice (Fricker, 2007) in which technologically nonaugmented agents are epistemically discriminated against because some aspects of their nonaugmented experiences no longer exist in the collective imagination of the mainstream augmented public. Sandor et al. (2015) acknowledge and take seriously this concern regarding the possibility of an 'augmented digital divide': "True AR presents a peculiar characteristic: it has the power to produce a radical divide of humanity into those who will live in an augmented world and others who will dwell in a world much less rich in information. This divide will affect the opportunities of individuals more profoundly than any previous technology, as it can continuously affect the perception of our immediate physical environment" (10).

One possible response to the digital divide is to argue for a human right to free internet access. Merten Regltiz, for example, argues that internet access qualifies as a human right in the digital age in virtue of being a necessary condition for the realization of democracy and the protection of other rights like free expression and freedom of assembly. According to Reglitz, the internet "is not a mere efficiency-enhancing technology but a medium for transforming human existence in an unprecedented way, which (as the UN puts it) 'by vastly expanding the capacities of individuals [...] contributes to the progress of humankind as a whole'" (Reglitz, 2020: 1). Following this line of thought, one might contend that in the case of the RWW, there is a *human right to augmentation*, meaning that agents should be supplied with unrestricted access to the RWW free of charge. Guaranteeing such a right through law could help mitigate the issues of epistemic injustice surrounding the digital divide.

7 Conclusion

This paper has examined some epistemological issues raised by emerging AR technology. Drawing on a range of relevant philosophical and empirical research, I showed how a near-future AR device called the Real-World Web threatens to exacerbate the problems of digital distraction, digital deception, and digital divergence. The RWW is poised to present new versions of these problems in the form of what I call the augmented attention economy, augmented skepticism, and the problem of other augmented minds. It is easy to write off these epistemic problems as being far-fetched or at least so far into the realm of science fiction that they are not worth taking into serious consideration now. I think this is a foolhardy attitude, especially in a world of unprecedented technological growth. The history of technology shows that sensible regulation and design implementation often lag behind technological innovation, meaning that our response to problems posed by new technologies is often reactive instead of proactive. It is typically only after a technology becomes widespread and starts to wreak havoc on society that we take serious measures to alleviate issues related to the technology. Correctly diagnosing and understanding different ethical, practical, and epistemological problems associated with emerging technologies enables us to devise policy-based and design-based solutions, or at least mitigation strategies, in a timely fashion. Throughout this paper, I proposed various mitigation strategies to the epistemic threats posed by the RWW, such as the idea of augmented ad blocking for the problem of digital distraction (Sect. 4), fluorescent (or see-through) design strategies for the problem of digital deception (Sect. 5), and the notion of a universal right to augmentation for the augmented digital divide concern (Sect. 6). In many respects, however, the discussion here has only scratched the surface of relevant epistemological and practical issues related to AR systems. I submit that the topic of augmented reality deserves more attention in the philosophy of technology given the increasing popularity and sophistication of AR devices, as well as how drastically the RWW promises to alter how agents interface with digital information.

Acknowledgements I am deeply grateful to Susan Schneider and two anonymous referees for offering helpful comments on previous versions of this paper. I would also like to thank Michael Lynch, Heather Battaly, and William Lycan for their mentorship as well as the Philosophy Department and Computer Science and Engineering Department at the University of Connecticut for their support over the years.

Declarations

Conflict of Interest The author declares no competing interests.

References

Agusanto, K., Li, L., Chuangui, Z., & Sing, N. W. (2003). Photorealistic rendering for augmented reality using environment illumination. *The second IEEE and ACM international symposium on mixed and augmented reality, Tokyo, Japan, October 10, 2003 (Piscataway, New Jersey* (pp. 208–216). IEEE).

Azuma, R. T. (1997). 'A survey of augmented reality.' *Presence: Teleoperators and Virtual Environments, 6*(4), 355–85.

Billinghurst, M. (2021). Grand challenges for augmented reality. *Frontiers in Virtual Reality*, 1–4.

Billinghurst, M., & Kato, H. (2002). Collaborative augmented reality. *Communications of the ACM, 45*(7), 64–70. https://doi.org/10.1145/514236.514265
Billinghurst, M., Kato, H., & Poupyrev, I. (2008). Tangible augmented reality. *ACM SIGGRAPH Asia, 7*(2), 1–10. https://doi.org/10.1145/1508044.1508051
Breen, D. E., Whitaker, R. T., Rose, E., & Tuceryan, M. (1996). Interactive occlusion and automatic object placement for augmented reality. *Computer Graphics Forum, 15.*
Bostanci, G. E., Kanwal, N., Ehsan, S., & Clark, A. (2013). User tracking methods for augmented reality. *International Journal of Human-Computer Studies., 5*, 93–98.
Bozdag, E., & van den Hoven, J. (2015). Breaking the filter bubble: Democracy and design. *Ethics and Information Technology, 17*(4), 249–265.
Brinkman, B. (2014). Ethics and pervasive augmented reality: Some challenges and approaches. In K. D. Pimple (Ed.), *Emerging pervasive information and communication technologies (PICT): Ethical challenges, opportunities and safeguards* (pp. 149–175). Springer.
Castro, C., & Pham, A. (2020). Is the attention economy noxious? *Philosophers' Imprint, 20*(17), 1–13.
Caudell, Thomas P., and David W. Mizell (1992). 'Augmented reality: An application of heads-up display technology to manual manufacturing processes.' *Proceedings of the 25th Hawaii International Conference on System Sciences*, 659–69.
Chalmers, D. J. (2017). The virtual and the real. *Disputatio, 9*(46), 309–352.
Chalmers, D. J. (2019). The virtual as the digital. *Disputatio, 11*(55), 453–486.
Chalmers, David J. (forthcoming). *Reality+: Virtual worlds and the problems of philosophy.* W.W. Norton and Penguin Press.
Chalmers, David J. (1996). *The conscious mind: In search of a fundamental theory*. Oxford University Press.
Clark, A., & Chalmers, D. J. (1998). The extended mind. *Analysis, 58*(1), 7–19.
Clowes, R. (2015). Thinking in the cloud: The cognitive incorporation of cloud-based technology. *Philosophy and Technology, 28*(2), 261–296.
Czarnecki, Tadeusz & Czarnecki, Bolesław (2017). 'Is augmented reality a source of new types of knowledge?' In José María Ariso (ed.), *Augmented Reality: Reflections on its Contribution to Knowledge Formation*. De Gruyter. pp. 151–170.
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. https://doi.org/10.1145/2556288.2557352
Dickey, R. M., Srikishen, N., Lipshultz, L. I., Spiess, P. E., Carrion, R. E., & Hakky, T. S. (2016). Augmented reality assisted surgery: A urologic training tool. *Asian Journal of Andrology, 18*(5), 732–734. https://doi.org/10.4103/1008-682X.166436
Eyal, N. (2014). *Hooked: How to Build Habit-Forming Products*. Portfolio Penguin.
Fallis, Don (forthcoming). 'The epistemic threat of deepfakes.' *Philosophy and Technology*:1–21.
Feng, S., Caire, R., Cortazar, B., Turan, M., Wong, A., & Ozcan, A. (2014). Immunochromatographic diagnostic test analysis using Google Glass. *ACS Nano, 8*(3), 3069–3079. https://doi.org/10.1021/nn500614k
Flach, J., & Holden, J. (1998). The reality of experience: Gibson's way. *Presence., 7*, 90–95.
Fricker, Miranda (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.
Friedman, B., & Kahn, P. H. (2000). New directions: A value-sensitive design approach to augmented reality. *DARE 2000: Design of Augmented Reality Environments, Elsinore, Denmark* (pp. 163–164). ACM.
Gibson, J. J. (1966). The senses considered as perceptual systems. *Synthese, 17*(2), 230–232.
Gibson, James J. (1979). *The ecological approach to visual perception.* Houghton Mifflin.
Grabarczyk, P., & Pokropski, M. (2016). Perception of affordances and experience of presence in virtual reality. *Avant: Trends in Interdisciplinary Studies, 7*(2), 25–44.
Hanin, Mark L. (2020). 'Theorizing digital distraction.' *Philosophy and Technology*:1–12.
Harris, T. (2018). 'The need for a new design ethics.' http://www.tristanharris.com/the-need- *for-a-new-design-ethics/*.
Kim, Hojoong, Kwon, Young-Tae, Lim, Hyo-Ryoung, Kim, Jong-Hoon, Kim, Yun-Soung, & Yeo, Woon-Hong. (2020). 'Recent advances in wearable sensors and integrated functional devices for virtual and augmented reality applications.' *Advanced Functional Materials*. Retrieved from https://par.nsf.gov/biblio/10235094. https://doi.org/10.1002/adfm.202005692

Lanier, J. (2018). *Ten arguments for deleting your social media accounts right now*. Henry Holt and Company.
Liberati, N. (2016). Augmented reality and ubiquitous computing: The hidden potentialities of augmented reality. *AI and Society, 31*(1), 17–28.
Liberati, N. (2018). Phenomenology, Pokémon Go, and other augmented reality games: A study of a life among digital objects. *Human Studies, 41*(2), 211–232.
Lingley, A. R., Ali, M., Liao, Y., Mirjalili, R., Klonner, M., Sopanen, M., Suihkonen, S., Shen, T., Otis, B. P., Lipsanen, H., & Parviz, B. A. (2011). A single-pixel wireless contact lens display. *Journal of Micromechanics and Microengineering, 21*(12), 125014–125021.
Liu, H., Zhang, G., and Bao, H. (2016). 'Robust keyframe-based monocular SLAM for augmented reality,' in *IEEE international symposium on mixed and augmented reality (ISMAR)*, Merida, Mexico, September 19–23, 2016 (IEEE), 1–10.
Matsumoto, T., Hashimoto, S., & Okude, N. (2008). The embodied web: Embodied web-services interaction with an umbrella for augmented city experiences. *Computer Animation and Virtual Worlds, 19*(1), 49–66.
McClelland, T. (2019). Representing our options: The perception of affordance for bodily and mental action. *Journal of Consciousness Studies, 26*(3–4), 155–180.
McClelland, Tom (2020) The mental affordance hypothesis. *Mind*.
McDonnell, N., & Wildman, N. (2019). Virtual reality: Digital or fictional? *Disputatio, 11*(55), 371–397.
Narumi, T., Nishizaka, S., Kajinami, T., Tanikawa, T., & Hirose, M. (2011). *Augmented reality flavors: Gustatory display based on edible marker and cross-modal interaction*. Paper presented at the Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Vancouver, BC, Canada.
Nee, A. Y., Ong, S. K., Chryssolouris, G., & Mourtzis, D. (2012). Augmented reality applications in design and manufacturing. *CIRP Annals, 61*(2), 657–679.
Neely, E. L. (2019). Augmented reality, augmented ethics: Who has the right to augment a particular physical space? *Ethics and Information Technology, 21*(1), 11–18.
Nguyen, C. Thi (2020). 'Echo chambers and epistemic bubbles.' *Episteme*:1–21.
NintendoSoup (2019). 'Pokemon Go officially hits 1 billion downloads worldwide.' Available at: https://nintendosoup.com/pokemon-go-officially-hits-1-billion-downloads-worldwide
Ophir, E., Nass, C., & Wagner, A. D. (2009). Cognitive control in media multitaskers. *Proceedings of the National Academy of Sciences, 106*(37), 15583–15587.
Orts-Escolano, S., Rhemann, C., Fanello, S., Chang, W., Kowdle, A., Degtyarev, Y., and Tankovich, V. (2016). 'Holoportation: Virtual 3D teleportation in real- time,' in *Proceedings of the 29th annual symposium on user interface software and technology*, Tokyo, Japan, October 2016 (New York, NY: Association for Computing Machinery), 741–754.
Palermos, Spyridon Orestis (2017). 'Augmented skepticism: The epistemological design of augmented reality.' In José María Ariso (ed.), *Augmented reality: Reflections on its contribution to knowledge formation*. De Gruyter. pp. 133–150.
Pariser, E. (2011). *The filter bubble: What the Internet is hiding from you*. Viking/Penguin Press.
Parviainen, Jaana (2017). 'Imagine never not knowing: An epistemological framework for understanding negative knowledge in augmented reality.' In José María Ariso (ed.), *Augmented reality: Reflections on its contribution to knowledge formation*. De Gruyter. pp. 195–216.
Pase, S. (2012). 'Ethical considerations in augmented reality applications,' in *Proceedings of the international conference on e-learninge-business, enterprise information systems, and e-Government EEE*. (The Steering Committee of the World Congress in Computer Science, Computer Engineering and Applied Computing (WorldComp)), 1. Las Vegas, Nevada, USA, July 16th - 19th, 2012.
Pezent, E., et al. (2019). Tasbi: Multisensory squeeze and vibrotactile wrist haptics for augmented and virtual reality. *IEEE World Haptics Conference (WHC), 2019*, 1–6. https://doi.org/10.1109/WHC.2019.881609
Regia-Corte, T., Marchal, M., Cirio, G., & Lécuyer, A. (2012). Perceiving affordances in virtual reality: influence of person and environmental properties in perception of standing on virtual grounds. *Virtual Reality, 17,* 17–28.
Reglitz, M. (2020). The human right to free Internet access. *Journal of Applied Philosophy, 37*(2), 314–331.
Roesner, F., Kohno, T., & Molnar, D. (2014). Security and privacy for augmented reality systems. *Communications of the ACM, 57*(4), 88–96.
Sandor, C., Fuchs, M., Cassinelli, A., Li, H., Newcombe, R., Yamamoto, G., et al. (2015). 'Breaking the barriers to true augmented reality.'

Santos, A. L. D, Lemos, D., Lindoso, J. E. F. and Teichrieb, V. (2012). 'Real time ray tracing for augmented reality,' *2012 14th Symposium on Virtual and Augmented Reality*, 2012, pp. 131–140, https://doi.org/10.1109/SVR.2012.8.

Sharp, T., Keskin, C., Robertson, D., Taylor, J., Shotton, J., Kim, D., et al. (2015). 'Accurate, robust, and flexible real-time hand tracking,' in *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, South Korea, April 2015 (New York, NY: Association for Computing Machinery), 3633–3642.

Shoemaker, S. (1982). The inverted spectrum. *Journal of Philosophy, 79*(July), 357–381.

Siegel, S. (2014). Affordances and the contents of perception. In B. Brogaard (Ed.), *Does perception have content?* (pp. 39–76). OUP.

Simonetta, G. (2015). 'The realism and ecology of augmented reality.' *Techné: Research in Philosophy and Technology, 19*(1), 92–112.

Smart, P. R. (2012). The web-extended mind. *Metaphilosophy, 43*(4), 446–463.

Smart, P. (2017). Extended cognition and the Internet: A review of current issues and controversies. *Philosophy and Technology, 30*(3), 357–390.

Smart, P., & Clowes, R. (2021). Intellectual virtues and Internet-extended knowledge. *Social Epistemology Review and Reply Collective, 10*(1), 7–21.

Smart, Paul R. (2013) 'Embodiment, cognition and the World Wide Web.' In L. Shapiro (Ed.), *The Routledge handbook of embodied cognition. Routledge,* New York, New York, USA

Smart, Paul (2018). 'Emerging digital technologies: Implications for extended conceptions of cognition and knowledge.' In J. Adam Carter, Andy Clark, Jesper Kallestrup, Spyridon Orestis Palermos & Duncan Pritchard (eds.), *Extended epistemology.* Oxford, UK: pp. 266–304.

Sugano, N., Kato, H., & Tachibana, K. (2003). The effects of shadow representation of virtual objects in augmented reality. *The second IEEE and ACM international symposium on mixed and augmented reality, Tokyo, Japan, October 10, 2003 (Piscataway, New Jersey* (pp. 76–83). IEEE).

Sullivan, S., & Laura & Reiner, Peter,. (2021). Digital wellness and persuasive technologies. *Philosophy and Technology, 34*(3), 413–424.

Sutherland, I. (1965). The ultimate display. *Proc. IFIP Congress, 2*, 506–508.

van Dijk, J. A. G. M. (2006). Digital divide research, achievements and shortcomings. *Poetics, 34*, 221–235.

Vold, K., & Hernandez-Orallo, J. (forthcoming). AI extenders and the ethics of mental health. In Ienca, M., & Jotterand, F. (Eds.), *Artificial intelligence in brain and mental health: Philosophical, ethical & policy issues*. Springer International Publishing.

Ward, A. F., Duke, K., Gneezy, A., & Bos, M. W. (2017). Brain drain: The mere presence of one's own smartphone reduces available cognitive capacity. *Journal of the Association for Consumer Research, 2*, 140–154.

Wassom, B. D. (2015). *Augmented reality law, privacy, and ethics: Law, society, and emerging AR technologies*. Elsevier.

Watson, Jamie Carlin. (2015). 'Filter bubbles and the public use of reason: Applying epistemology to the newsfeed.' In *Social epistemology and technology: Toward public self- awareness regarding technological mediation.* Ed. Frank Scalambrino. London: Rowman & Littlefield.

Weiser, M. (1991). 'The computer for the 21st century.' *In: Scientific American, 265*(3), 94–104.

Weiser, M., & Brown, J. S. (1996). Designing calm technology. *PowerGrid Journal, 1*(1), 75–85.

Wellner, G. (2013). No longer a phone: The cellphone as an enabler of augmented reality. *Transfers, 3*(2), 70–88.

Williams, James (2018). *Stand out of our light: Freedom and resistance in the attention economy*. Cambridge University Press.

Wolf, M. J., Grodzinsky, F., & Miller, K. (2015). Augmented reality all around us. *Acm Sigcas Computers and Society, 45*(3), 126–131.

Zambrano, A., & Pickard, C. (2018). A defense of ad blocking and consumer inattention. *Ethics and Information Technology, 20*(3), 143–155.

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Public Affairs.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.