To appear in *Augmented Reality. Reflections on its Contribution to Knowledge Formation*, José María Ariso (ed), 2017, De Gruyter (series: Berlin Studies in Knowledge Research), Berlin & New York.

# Augmented Skepticism:
## The Epistemological Design of Augmented Reality

S. Orestis Palermos
University *of* Edinburgh

**Abstract.** In order to solve the problem of the traditional account of knowledge, according to which justification is the ability to provide reflectively accessible positive reasons in support of one's beliefs, a number of epistemologists have suggested that knowledge is true belief that is the product of cognitive ability. According to this alternative, a belief-forming process may count as a knowledge-conducive cognitive ability if and only if it has been cognitively integrated on the basis of processes of mutual interactions with other aspects of the agents' cognitive system. One of the advantages of this approach is that it allows knowledge and justification to be extended to such artifacts as telescopes, microscopes, smartphones and augmented reality (AR) systems. AR systems, however, rely on deceptive reality augmentations that could significantly deteriorate the epistemic efficiency of users' cognitively integrated natures. This could lead to a form of 'augmented skepticism', whereby it will be impossible to tell augmented from physical reality apart. In order to solve this problem, epistemology should play an active role in the design of future AR systems and practices. To this end, this chapter puts forward some initial suggestions, concerning the training of AR users and the design of certain reality augmentation features, in order to ensure that everyday epistemic practices won't be disrupted by the introduction of emerging AR technologies.

**Keywords:** Augmented Reality; Extended Cognition; Epistemology; Philosophy of Cognitive Science; Cognitive Integration; Extended Knowledge; Philosophical Engineering; Augmented Skepticism

> Everything we see hides another thing, we always want to see what is hidden by what we see. There is an interest in that which is hidden and which the visible does not show us. This interest can take the form of a quite intense feeling, a sort of conflict, one might say, between the visible that is hidden and the visible that is present.
>
> (René Magritte on *The Son of Man*)

## Introduction

Weeks after the release of Pokémon Go, Police are offering safety advice to users of the popular online game and reminding players to concentrate on the real world when catching Pokémon. Car accidents, property trespassing, carelessly crossing the road, walking through landmines, and wandering in dangerous areas at inappropriate times of the day have raised a number of concerns, all related to the attention deficiency of overexcited users. Yet failing to concentrate on the real world is not the only and certainly not the most worrying aspect of augmented reality (AR).[1]

'Seeing is believing' could so far be hardly doubted in most ordinary contexts. Yet this fundamental aspect of our everyday epistemic life is likely to be soon under serious threat by the advent of AR. As AR will become ubiquitous, it will likely take over most aspects of our daily interactions with surrounding objects and human beings, making it practically impossible to distance ourselves from this added dimension of future society, much in the same way that most people can no more leave their house without making sure they have their mobile phones on them. There is, no doubt, a great potential in this emerging technology, which promises to enrich our lives beyond imagination. But its users may also be exposed to the serious danger of being unable to tell reality and augmented reality apart.

This form of future 'augmented scepticism' cannot be neglected and important steps need to be taken with regards to the design of future AR systems as well as teaching users how to employ the emerging technology in order to avoid this looming epistemic threat. By focusing on recent advances within contemporary epistemology and philosophy of mind and cognitive science, and especially the notion of *cognitive integration*, this chapter attempts to address this concern and provide advice that could secure our knowledge of the external world while also allowing our knowledge to be extended beyond our biological capacities, by taking advantage of the opportunities offered by AR.

## Knowledge and Cognitive Integration

The received epistemological view holds that knowledge is justified true belief. 'Justification', however, is a term of art that can be given a number of different interpretations. According to the traditional account of knowledge, justification is a form of ability to provide explicit positive reasons in support of one's beliefs by reflection alone.[2] This is a familiar demand. We are many times asked to provide explicit reasons in support of our epistemic statements as well as in support of our reasons for claiming that we know such statements and so on. Nevertheless, however common this practice may be, it cannot really represent a universal theory of knowledge and justification as it generates serious problems, both from a theoretical and a practical point of view.

From a theoretical perspective, demanding to always be in a position to offer reasons in support of one's beliefs by reflection alone has the paralyzing epistemic effect of disallowing all perceptual and empirical knowledge. Technically, asking for one to justify one's perceptual

[1] Strictly speaking, given how the current version of Pokémon Go works, it is a combination of both virtual and augmented reality.

[2] Within contemporary epistemology, this is known as epistemic internalism. For classical defenses of this view see Chisholm (1977) and BonJour (1985, Chap. 2). See also Steup (1999), Pryor (2001, p. 3), BonJour (2002), Pappas (2005), and Poston (2008).

and empirical beliefs by reflection alone poses the requirement that there be necessary support relations between one's empirical and perceptual beliefs and one's evidence for holding them. Hume's problem of induction demonstrates, however, that this is impossible.[3]

Similarly, from a practical perspective, there is a number of belief-forming processes, such as vision, hearing and memory, which are supposed to be knowledge-conducive, even though most epistemic agents have no idea how they work or why they are reliable.[4] Accordingly, when we acquire knowledge on their basis, it seems incorrect to require explicit positive reasons in their support.

To solve this long standing problem, several epistemologists have recently suggested that we should give up the aforementioned understanding of justification, and instead embrace a weaker alternative. According to this weaker alternative in order for one's true beliefs to qualify as knowledge, they must simply be the product of a belief-forming process that counts as a cognitive ability.[5] This is known as the *ability intuition on knowledge* and its intuitive appeal comes from the fact that cognitive abilities do seem to be the sort of belief-forming processes that can generate knowledge, even if one has no explicit positive reasons to offer in their support.[6] No one needs to explain why their vision or hearing is reliable when they come to acquire knowledge on their basis, after all.

If this is the way to approach knowledge and justification, however, two central questions need to be further addressed: (1) When does a process count as a cognitive ability and thereby as knowledge conducive, and—depending on how we answer (1)—(2) what is the sense in which one can be justified/epistemically responsible on the basis of one's cognitive abilities, but without requiring to offer any explicit reasons in their support?

In order to answer these two important questions, epistemologists have turned to the concept of *cognitive integration*. Recently, it has been proposed that in order for a process to count as a cognitive ability (and thereby as knowledge-conducive) it must have been *cognitively integrated*, where cognitive integration is a "function of cooperation and interaction, or

---

[3] The problem of induction is well known. We form our beliefs about unobserved matters of fact and the external world on the basis of evidence provided by past and present observations and sensory appearances, respectively. In order for the support relations between our empirical and perceptual beliefs and the evidence offered in their support to be necessary, we also need the further assumptions that the future will resemble the past and that sensory appearances are reliable indications to reality, respectively. The problem is that both of these assumptions rely for their support on what they assert. Consequently, given that circular reasoning is invalid, there are no necessary support relations between our empirical beliefs and the evidence offered in their support. Accordingly, the conclusion that has been traditionally drawn is that our empirical and perceptual beliefs cannot amount to knowledge. For more details on how to reconstruct Hume's skepticism along these lines, see (Greco 1999).

[4] This claim may generalize to all epistemic agents. There are no widely received or established views within cognitive science, regarding the mechanisms underlying any of the above belief-forming processes.

[5] Within contemporary epistemology, this is known as virtue reliabilism. There is a still weaker alternative conception of knowledge and justification within the literature known as process reliabilism. For an overview of process reliabilism see (Goldman and Beddor 2015). For a number of arguments against the view and why virtue reliabilism is to be preferred see (Greco 1999, 2010, Pritchard 2010; Palermos 2014b)

6 The idea that knowledge must be grounded in cognitive abilities can be traced back to the writings of (Sosa 1988, 1993) and Plantinga (1993*a*, 1993*b*). For more recent approaches to this intuition, see Greco (1999; 2004; 2007; 2010) and Pritchard (2009, 2010a, 2010b, 2010, 2012).

cooperative interaction with other aspects of the cognitive system" (Greco 2010, 152). Accordingly, the answer to the first question is that a process may count as a cognitive ability (and thereby as knowledge-conducive) so long as it has been cognitively integrated on the basis of processes of mutual interactions with other aspects of the cognitive system.

One of the virtues of this approach to knowledge and justification is that it is fairly straightforward: In order for a reliable belief-forming process to count as knowledge-conducive, it must also count as a cognitive ability, and, in order for that to be the case, the relevant belief-forming process must mutually interact with other aspects of the cognitive system. Yet an additional advantage of this approach is that it can also provide a satisfactory response to the second question we posed above—i.e., what is the specific sense in which one can be justified/epistemically responsible on the basis of one's cognitive abilities, even in the absence of any explicit reasons in support of their reliability? The key, again, is to focus on the cooperative and interconnected nature of cognitive abilities: If one's belief-forming process interacts cooperatively with other aspects of one's cognitive system then it can be continuously monitored in the background such that *if* there is something wrong with it, *then* the agent will be able to notice this and respond appropriately. Otherwise—if the agent has no negative beliefs about his/her belief-forming process—he/she can be subjectively justified/epistemically responsible in employing the relevant process *by default*, even if he/she has absolutely no positive beliefs as to whether or why it might be reliable.

For example, in order for agent *S* to responsibly hold the belief that there is man standing in front of her, *S* does not need to offer explicit, positive reasons in support of the reliability of her visual system. Instead, provided that *S*'s visual system is interconnected with the rest of her cognitive system, then, in the mere lack of defeaters against the reliability of her visual perception, *S* can take herself to be epistemically responsible in holding the relevant belief *by default*. Had her working memory alerted her to the fact that the lighting conditions were not good, had she felt extremely tired, had her long term memory reminded her that she is watching a magic show, or had she tried to touch the person without receiving the expected tactile feedback, she would refrain from accepting the visually formed belief, no matter how truth-like it would appear to her. Nevertheless, in the absence of any such negative reasons against her belief, she can take herself to be epistemically responsible in holding the automatically delivered visual belief, *by default* (Palermos, 2014b).

This way, we can make sense of the commonly held idiom that 'seeing is believing', or at least, 'seeing is believing, unless there are reasons to be believe it is not.'

## Extended Knowledge and Cognitive Integration

But is this always the case, or just when we perceive the world through our biological equipment? Recent studies at the intersection of epistemology and philosophy of mind and cognitive science indicate that knowledge and justification can be technologically extended. (Pritchard 2010c; Palermos 2011, 2014b, 2015, 2016; Palermos and Pritchard 2013; Carter, Kallestrup, Palermos and Pritchard 2014)

Over the last two decades, philosophy of mind and cognitive science has become increasingly receptive to the idea that cognition is not head-bound but instead potentially extended to the artifacts we mutually interact with. Broadly known as the current of active externalism, this idea has been expressed under a number of headings by several philosophers and cognitive scientists (Clark and Chalmers 1998; Rowlands 1999; Wilson 2000; Wilson 2004; Menary 2007). One of the most influential formulations—perhaps the most influential—is known as the hypothesis of extended cognition and it holds that "the actual local operations that realize certain forms of human cognizing include inextricable tangles of feedback, feedforward and feed-around loops: loops that promiscuously crisscross the boundaries of brain, body and world" (Clark 2007, sec. 2). A list of examples of interactive, cognition extending equipment would include telescopes, microscopes, GPS systems, even pen and paper when trying to solve complex scientific problems (Palermos 2015) or while performing simple multiplication tasks.

Think about a three-digit multiplication problem such as 987 times 789. It is true that few if any of us can solve this problem by looking at or contemplating on it. We may only perform the multiplication process by using pen and paper to externalize the problem in symbols. Then we can serially proceed to its solution by performing simpler multiplications, starting with 9 times 7, and externally storing the results of the process for use in later stages. The process involves eye-hand motor coordination and it is not simply performed within the head of the person reciting the times tables. It involves intricate, continuous interactions between brain, hand, pen and paper, all the while it is being transparently regulated by the normative aspects of the notational/representational system involved—for instance, that we cannot multiply by infinity, that we must write the next digit under the second to last digit of the number above, what operation we must perform next and so on.[7]

Proponents of the hypothesis of extended cognition note that in such cases we can talk of an extended cognitive system that consists of both biological and technological resources, because the completion of the relevant cognitive task (e.g., performing the multiplication task) involves non-linear, cooperative interactions between the two components. According to dynamical systems theory (DST)—i.e., the most promising mathematical framework for modeling such dynamically interacting systems—when this is the case we have to postulate an overall coupled system that consists of all the mutually interdependent components at the

[7] For the importance of the normative aspects of the external representational systems in explaining cognition see (Menary 2007).

same time.[8] According to a dynamical interpretation of the hypothesis of extended cognition, when two (or more) components mutually interact with each other in order to complete a cognitive task, they give rise to an extended cognitive system that consists of all of them at the same time.[9]

This brings to the fore the possibility that knowledge-conducive cognitive abilities can be extended to the artifacts we employ. This is because epistemology and philosophy of mind and cognitive science put forward the same condition in order for a process to count as *cognitively integrated*, and thereby knowledge-conducive: Just as philosophers of mind claim that a cognitive system is integrated when its contributing parts engage in ongoing reciprocal interactions (independently of *where* these parts may be located), so epistemologists claim that cognitive integration of a belief-forming process (be it internal or external to the agent's organism) is a matter of cooperative interactions with other parts of the cognitive system.[10] The theoretical wedding of the two disciplines suggests there is no reason to disallow the belief-forming processes of extended or even distributed cognitive systems from counting as knowledge-conducive.

Provided that the relevant system is cognitively integrated on the basis of the mutual interactions of its component parts, it can generate epistemically responsible/justified beliefs, independently of whether it is organism-bound or extended. The ongoing interactivity of its component parts—i.e., its cognitively integrated nature—allows the system to be in a position such that if there is anything wrong with the overall process of forming beliefs, the system will be alerted to it and respond appropriately. Otherwise, if there is nothing wrong, the system can accept the deliverances of its belief-forming processes by default, without the further requirement to provide explicit positive reasons in their support. This is a form of justification/epistemic responsibility that does not belong to any of the component parts but to the relevant system as a whole. The reason is that it does not arise on the basis of any component parts operating in isolation but instead on their ongoing interactivity, which, according to DST, belongs to the system as whole.

For example, it is possible to use the above approach in order to explain how a subject might come to perceive the world on the basis of a Tactile Visual Substitution System (TVSS), while also holding fast to the idea that knowledge is belief that is true in virtue of *cognitive ability* (i.e. the ability intuition on knowledge). A TVSS comprises of a mini video

---

[8] For more details behind this rationale and an extensive defense of this claim see (Palermos 2014a). For an introduction to DST, see (Abraham, Abraham and Shaw, 1990).

[9] By contrast, using a ladder to paint the ceiling, heating food with microwaves, supermarket lists, turning the lamp on to see in a dark room, etc. won't qualify as cases of cognitive extension because in such cases there is no ongoing mutual interactivity between the agent and the involved artifact.

[10] Elsewhere (Palermos 2011, Palermos 2014b), it has been argued that both disciplines also put forward the same broad, common sense functionalist intuitions on what is required from a process to count as a cognitive ability. Briefly, both views state that the process must be (a) normal and reliable, (b) one of the agent's habits/dispositions and (c) integrated into the rest of the agent's cognitive character/system.

camera attached on a pair of glasses, which converts the visual input into tactile stimulation under the agent's tongue or her forehead. By moving around and on the basis of the associated sensorimotor contingencies,[11] blind patients quickly start perceiving shapes and objects and orienting themselves in space. Occasionally, they also offer reports of feeling as if they are *seeing* objects, indicating that they are enjoying phenomenal qualities very close to those of the original sense modality that is being substituted. In light of DST, seeing through a TVSS qualifies as a case of cognitive extension, because it is a dynamical process that involves ongoing reciprocal interactions between the agent and the artifact. By moving around, the agent affects the input of the mini-video camera, which continuously affects the tactile stimulation she will receive on her tongue or forehead by the TVSS, which then continuously affects how she will move around and so on. Eventually, as the process unfolds, the coupled system of *the agent and her TVSS* is able to identify—that is, see—shapes and objects in space.

## Augmented Skepticism

Given the way augmented reality systems work, they have the potential to qualify as *cognitively integrated* and thereby knowledge-conducive extensions of biological cognition.

Most modern augmented reality systems combine the input from hardware components such as digital cameras, accelerometers, global positioning systems (GPS), gyroscopes, solid state compasses, and wireless sensors with simultaneous localization and mapping (SLAM) software, in order to track the position and orientation of the user's head and overlay computer data and graphics to her visual field in real time. By moving around with the AR system, the user affects the input received by the hardware components, which continuously feeds in to the SLAM software. In turn, the SLAM software keeps constructing and updating a map of the user's unknown environment while simultaneously keeping track of the user's position in the physical world, the way she is pointing the device at and the axis the device is operating in. This constant interplay between the user, the AR hardware and the AR software allows the system to display computer-generated images on the user's field of perception and allows the user to visually interact with these virtual images while she moves in space as if they were real, physical objects.

In light of epistemology and philosophy of mind and cognitive science, this advanced degree of ongoing mutual interactivity between the user and the AR system indicates that AR

[11] For a recent review on TVSS, see Bach-y-Rita and Kercel (2003). For a full account of how sensorimotor knowledge is constitutive of perception see (Noë 2004). "The basic claim of the enactive approach is that the perceiver's ability to perceive is constituted (in part) by sensorimotor knowledge (i.e. by practical grasp of the way sensory stimulation varies as the perceiver moves)." (Noë 2004, 12) "What the perception is, however, is not a process in the brain, but a kind of skillful activity on the part of the animal as a whole". (Noë 2004, 2). "Perception is not something that happens to us or in us, it is something we do". (Noë 2004, 1). Sensorimotor dependencies are relations between movements or change and sensory stimulation. It is the practical knowledge of loops relating external objects and their properties with recurring patterns of change in sensory stimulation. These patterns of change may be caused by the moving subject, the moving object, the ambient environment (changes in illumination) and so on.

can become a powerful technology for extending our knowledge beyond the epistemic abilities provided by our organismic cognitive capacities. A number of emerging applications across a multitude of disciplines indicate this clearly.

Users can perceive electromagnetic radio waves overlaid in exact alignment with their actual position in space. AR can also be used to assist archaeological research, by superimposing archaeological features onto modern landscapes, allowing archaeologists to draw inferences about site placement and configuration. AR archaeology applications can assist users reconstruct ruins, buildings and landscapes as they formerly existed. Architects and civil engineers can employ the technology to visualize future building projects. Computer-generated images of buildings can be overlaid into a real life local view of a property before the construction process begins. Architecture sight-seeing can be enhanced with AR applications allowing users to virtually see through the walls of buildings and gain access to visual information about interior objects and layout. With recent improvements to GPS accuracy, construction companies are able to use augmented reality to visualize georeferenced models of construction sites, underground structures, cables and pipes.

Similarly, there is a number of potential commercial uses. AR can enhance product previews such as allowing consumers to view what's inside a product's packaging without opening it. It can also be used in order to facilitate the selection of products from a catalogue or a kiosk. AR users could gain access to additional content such as customization options and images or videos of the product in its use. Such technologies are already in use. It is possible, for example, to design printed marketing material so that it can bear certain "trigger" images that, when scanned by an AR device, they activate a video version of the promotional material.

AR can also make significant contributions to health and safety. Imagine a rescue pilot who is looking for a lost hiker in a forest. Augmented reality systems can provide geographic awareness of forest road names and locations. As a result, the rescuer can more easily detect the hiker knowing the geographic context provided by the AR system. Similarly, AR can be used to let a surgeon look inside a patient by combining one source of images such as an f-MRI scan with another such as video.

AR can also augment the effectiveness of navigation devices. Directions can be displayed on a car’s windshield, while also indicating weather, terrain, road conditions and traffic information as well as alerts to potential hazards. Augmented reality applications can enhance a user's travel experience by providing real time informational displays of her location and its features, as well as access to comments of previous visitors of the site. AR applications can allow archaeological site visitors to experience simulations of historical events, places and objects by overlaying them into their view of a landscape. They can also offer location information by audio, calling attention to features of interest as they become visible to the user.

The above examples make it obvious that AR has the potential to permeate and enrich our everyday lives in a variety of ways. As AR technologies become less intrusive and more transparent, moving from hand held devices, to AR glasses and finally to contact lenses, AR will possibly not only penetrate every aspect of our lives but will become a constant, additional layer to physical reality that users will be practically unable to disengage from. Short films *Sight* (https://vimeo.com/46304267) and *Hyper-Reality* (https://vimeo.com/166807261) provide good tasters of how the augmented future might soon look like.

AR therefore promises to provide a great opportunity for extending our knowledge in a variety of new and exciting ways. At the same time, however, it also poses the serious threat of obstructing our knowledge of the external world. Contrary to other forms of extended cognitive systems, AR is specifically designed to generate and operate on the basis of unreal yet deceivingly truth-like mimicries of the external world in a way that users won't be able to distinguish augmented images from actual images of the world.

Of course, the integrated nature of our cognitive systems may still be in a good position to single out reality augmentations that cannot be easily confused as parts of physical reality. For example, floating prize tags above products or fluorescent navigation arrows in our visual field won't be of particular concern. On the basis of cognitive integration, our previous experience and knowledge of the external world will allow us to perceive such items as reality augmentations. Other aspects of augmented experience, however, are going to be troubling.

Consider, for example, *S*'s mundane experience of visually perceiving that a person is standing opposite her. *S* will be considerably worse off holding such a belief in an epistemically responsible manner while having her AR system turned on than when she has it turned off. The possibility of having real-like yet virtual representations being superimposed on one's perception of the physical world will require a much more thorough background check by *S*'s integrated cognitive system before she can believe what she perceives. Normally, the presence of good-lighting and a relatively stable experience, along with the absence of any beliefs regarding the possibility of being tricked by a magician or undergoing drug-induced hallucinations, would be more than enough for *S* to know that there is a person standing opposite her. An AR experience, however, would essentially amount to participating in a magic show. As such, believing what one sees would additionally require making haptic checks or being sensitive to additional cues that could potentially warn *S*' cognitively integrated nature to the fact that she is in a context where the presence of AR avatars is to be expected.

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.[12]

## Future Use and Design

AR therefore has the potential to both extend and distract our organismic epistemic capacities. Of course, technology optimists may disregard the above worries as being exaggerated. One could turn their AR systems off anytime they liked, thereby eliminating the threat of 'augmented skepticism' at the push of a button. But how realistic is such optimism?

Considering the present-day analogue of owning a smart-phone, how often do we turn them off? Mobile phones are significantly less intrusive and attention-grabbing than future augmented reality technologies such as AR glasses and AR lenses are going to be. Smart-phones require their users to actively look at the screen instead of having information automatically pushed within their visual field. Yet mobile phone addiction has already started posing real life threats:

> In the case of cell-phones, such an addiction may begin when an initially benign behavior with little or no harmful consequences—such as owning a cell-phone for safety purposes—begins to evoke negative consequences and the user becomes increasingly dependent upon its use. Owning a cell-phone for purposes of safety, for instance, eventually becomes secondary to sending and receiving text messages or visiting online social networking sites; eventually, the cell-phone user may engage in increasingly dangerous behaviors such as texting while driving. Ultimately, the cell-phone user reaches a "tipping point" where he/she can no longer control their cell-phone use or the negative consequences from its over-use (Roberts, Yaya and Manolis 2014, 255).

Responsible theorizing and future planning and design cannot therefore rest on unsubstantiated optimism, especially when relevant evidence points in the opposite direction. Future AR technologies are more likely than not to storm users' visual fields with push-on notifications, advertisements, personalized suggestions and reminders. Such reality augmentations could, in the best-case scenario, obstruct the user's perception of the external world and, in the worst-case scenario, cause severe disorientation with regards to what may be part of actual reality.

Careful planning and design, however, can reduce or even eliminate such risks. The preceding epistemological remarks on the role of cognitive integration can offer significant guidance to this end. Previously we noted that epistemic responsibility and justification rely on

[12] Indeed, it could have a destructive effect. One of the most promising ways to avoid the threat of radical skepticism is to note that our everyday beliefs are modally *safe*. However, this strategy works only on the assumption that radical skeptical hypotheses are modally far off from the actual world (see for example, Pritchard 2013, §2). Practically speaking, this assumption has so far been easy to grant. Nevertheless, the advent of AR technologies could make radically skeptical scenarios modally close, thereby seriously questioning our psychological dismissal of radical skepticism on the basis of practical considerations.

the mutual interactivity of the agent's belief-forming processes. If there is something wrong with the way the agent is currently forming her beliefs, then it will clash with at least one of the agent's belief-forming processes running in the background, such that the agent will take notice and respond appropriately. Otherwise, if there is nothing wrong, the agent can accept the deliverances of her belief-forming process *by default*.

Given that AR overlays augmentations on one's visual field, many of which might be deceptively real, one initial suggestion is to attempt to teach users how to employ the technology in a way that can diminish the ensuing 'augmented skepticism'. While it is difficult to imagine how future AR will actually look like, a generic solution to this problem may include the progressive training of AR users to recognize and automatically be aware of settings and social contexts in which deceptive reality augmentations are likely to be present. In such cases, users will have to be aware that relying on what they perceive won't be safe. Instead, they will need to employ their cognitively integrated nature more than it is normally required by performing additional background checks that will involve supplementary interactions with the perceived item (e.g., reaching out for the item in order to test whether it will provide the corresponding haptic feedback).

Key to the above solution is that users will be able to tell deceptive reality augmentations from non-deceptive ones apart. It assumes that even though users may be tricked by reality augmentations that look like deceptive representations of physical reality, they can easily spot augmentations that are unlikely to be found in physical reality (e.g., floating price tags above products, or navigational arrows pointing users in the right direction). This ability of our cognitively integrated natures relies on extensive previous experience of interacting with the physical world.

But what happens if the user has never had the opportunity to become thoroughly acquainted with the physical world outside AR? Given how attractive digital technologies are to children, this is a developmental danger that future educational systems and upbringing must take into consideration. It may well sound as yet another exaggerated threat, but given the potential prevalence of AR in future societies, it may not be easily disregarded as far-fetched. Should that ever become the case, children and students should be encouraged to spend as much of their day interacting with the actual physical world alone, or they may fail to enhance their cognitively integrated nature with the expectations that will be required to tell most instances of augmented and physical reality apart—even if reality augmentations are specifically designed to stand out from physical reality.

Future AR users should therefore prime their cognitively integrated nature to identify non-deceptive augmentations as well as the contexts and settings in which deceptive augmentations are likely to appear. Yet despite such measures, users' epistemic standing may still be severely compromised. Not at all unlikely, the contexts and settings in which deceptive augmentations may appear could be widespread or even ubiquitous. If that turns out to be the case, users' ability to perceive the external world would be severely limited and slowed down, due to having to perform a number of additional—presently unnecessary—background checks with every step they'd take. Eventually, their experience would amount to walking through a mirror room.

A solution to this problem would require turning our attention away from the users' practices and towards the design of AR. AR developers would have to make sure that all augmentations bear features that would allow them to clearly and immediately stand out from the physical elements in the world without the need of unrealistically burdensome checks on the part of the users. The design of future AR systems should not pose unrealistic demands on the users' cognitively integrated nature. Reality augmentations should automatically stand out as such, leaving minimal room for confusion or misinterpretation. For example, they should be delineated with fluorescent borders, have a see-through effect or both. In fact, to ensure users' epistemic ease and safety, such AR design specifications could even be enforced via public policies and the law.[13]

Instead, a completely immersive experience, where virtual images could be entirely indistinguishable from physical reality could be retained for virtual reality, where the user's awareness of her physical disengagement will allow her to fully and safely enjoy the experience of mediated reality.

## Conclusion

In order to solve the problem of the traditional account of knowledge, according to which justification is the ability to provide reflectively accessible positive reasons in support of one's beliefs, a number of epistemologists have suggested that knowledge is true belief that is the product of a cognitive ability. According to this alternative, a belief-forming process may count as a knowledge-conducive cognitive ability if and only if it has been cognitively integrated on the basis of processes of mutual interactions with other aspects of the agents' cognitive system. One of the advantages of this approach is that it allows knowledge and justification to be extended to such artifacts as telescopes, microscopes, smartphones and AR systems. AR systems, however, rely on deceptive reality augmentations that could significantly deteriorate the epistemic efficiency of users' cognitively integrated natures. This could lead to a form of 'augmented skepticism', whereby it will be impossible to tell augmented from physical reality apart. In order to solve this problem, epistemology should play an active role in the design of future AR systems and practices. To this end, this chapter has put forward

[13] For further considerations on how the hypothesis of extended cognition might invite a reconceptualisation of current legal theorising and practices, and especially of how we should perceive the right against personal assault, see (Carter and Palermos, forthcoming).

some initial suggestions, concerning the training of AR users and the design of certain reality augmentation features. This is but a first step to ensuring that our everyday epistemic practices won't be easily disrupted by the advent of AR technologies. To avoid such and similar threats it is important to not undermine the input that philosophical engineering (Halpin 2013; Hendler & Berners-Lee, 2010; Halpin et al. 2010; Palermos forthcoming), in general, and epistemological design, in particular, can provide to the development of emerging and future technologies.

## Acknowledgements

I am very thankful to Duncan Pritchard and Panagiotis Kourtesis for their thoughtful comments on a previous draft of this chapter.

## References

Abraham, F. S., Abraham, R. H., & Shaw, C. (1990). A Visual Introduction to Dynamical Systems Theory for Psychology. Santa Cruz, CA: Aerial Pr.

Bach-y-Rita, P., and S. W. Kercel. (2003). 'Sensory substitution and the human- machine interface', *Trends in Cognitive Science* 7, no.12 541-6.

BonJour, L. (1985). *The Structure of Empirical Knowledge*. Harvard University Press.

—— (2002). 'Internalism and Externalism', *Oxford Handbook of Epistemology*, (ed.) P. Moser, 234-64, Oxford: Oxford University Press.

Carter, J. A., Kallestrup, J., Palermos, S. O., & Pritchard, D. (2014). Varieties of externalism. *Philosophical Issues*, *24*(1), 63-109.

Carter, J. A., and Palermos, S. O. (forthcoming). The Ethics of Extended Cognition: Is Having Your Computer Compromised a Personal Assault? *The Journal of American Philosophical Association.*

Chisholm, R. M. (1977). *Theory of knowledge*. Prentice-Hall.

Clark, A. (2007). 'Curing Cognitive Hiccups: A Defence of the Extended Mind', *The Journal of Philosophy,* 104: 163-192.

Clark, A., & Chalmers, D. (1998). The extended mind. *Ananlysis*, 58(1), 7–19.

Goldman, Alvin and Beddor, Bob, "Reliabilist Epistemology", *The Stanford Encyclopedia of Philosophy* (Winter 2015 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2015/entries/reliabilism/>.

Greco, J. (1999). 'Agent Reliabilism', in *Philosophical Perspectives 13:* Epistemology (1999). James Tomberlin (ed.), Atascadero, CA: Ridgeview Press, pp. 273-296.

—— (2004). 'Knowledge As Credit For True Belief', in *Intellectual Virtue: Perspectives from Ethics and Epistemology*. M. DePaul & L. Zagzebski (eds.), Oxford: Oxford University Press.

—— (2007) 'The Nature of Ability and the Purpose of Knowledge', *Philosophical Issues* 17, pp.

57- 69.

—— (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity.* Cambridge University Press.

Halpin, H., Clark, A., & Wheeler, M. (2010). Towards a philosophy of the web: representation, enaction, collective intelligence. In: *Proceedings of the WebSci10: Extending the Frontiers of Society On-Line*.

Halpin, H. (2013). Does the web extend the mind?. In *Proceedings of the 5th annual ACM web science conference* (pp. 139-147). ACM.

Hendler, J., & Berners-Lee, T. (2010). From the Semantic Web to social machines: A research challenge for AI on the World Wide Web. *Artificial Intelligence*, *174*(2), 156-161.

Menary, R. (2007). *Cognitive integration: Mind and cognition unbound*. Basingstoke: Palgrave McMillan.

Noë, A. (2004). *Action in Perception.* Cambridge, MA:MIT Press.

Palermos, O. (2011). 'Belief-Forming Processes, Extended', *Review of Philosophy and Psychology* 2, 741-65.

Palermos, S. O. (2014a). 'Loops, Constitution, and Cognitive Extension', *Cognitive Systems Research,* 27: 25‑41.

Palermos, S. O. (2014b). Knowledge and cognitive integration. *Synthese,191(8), 1931-1951.*

Palermos, S. O. (2015). Active externalism, virtue reliabilism and scientific knowledge. *Synthese*, *192*(9), 2955-2986.

Palermos, S. O. (2016). Spreading the Credit: Virtue Reliabilism and Weak Epistemic Anti-Individualism. *Erkenntnis*, *81*(2), 305-334.

Palermos. S. O. (forthcoming). Social Machines: A Philosophical Engineering. *Phenomenology and the Cognitive Sciences.*

Palermos, S. O., *&* Pritchard, D. (2013*)*. 'Extended Knowledge and Social Epistemology', *Social Epistemology Review and Reply Collective* 2 (8): 105-120.

Pappas, G. (2005). 'Internalist *versus* Externalist Conceptions of Epistemic Justification', *Stanford Encyclopaedia of Philosophy*, (ed.) E. Zalta, http://plato.stanford.edu/entries/justep-intext/

Poston, T. (2008). 'Internalism and Externalism in Epistemology', *Internet Encyclopaedia of Philosophy*, (eds.) B. Dowden *&* J. Fieser, www.iep.utm.edu/int-ext/

Pritchard, D. (*forthcoming*). 'Anti-Luck Virtue Epistemology', *Journal of Philosophy.*

—— (2009). *Knowledge,* London: Palgrave Macmillan.

—— (2010*a*). 'Knowledge and Understanding', in A. Haddock, A. Millar & D. H. Pritchard, *The Nature and Value of Knowledge: Three Investigations*, Oxford: Oxford University Press.

—— (2010*b*). 'Cognitive Ability and the Extended Cognition Thesis'. *Synthese.*

—— (2013). *What is this thing called knowledge?*. Routledge.

Pryor, J. (2001). 'Highlights of Recent Epistemology', *British Journal for the Philosophy of Science* 52, 95-124.

Roberts, J., Yaya, L., & Manolis, C. (2014). The invisible addiction: Cell-phone activities and addiction among male and female college students. *Journal of behavioral addictions*, *3*(4), 254-265.

Rowlands, M. (2009). 'Extended Cognition and the Mark of the Cognitive'. *Philosophical Psychology,* 22(1); 1-19.

Steup, M. (1999). 'A Defense of Internalism', *The Theory of Knowledge: Classical and Contemporary Readings* (3rd Ed.), (ed.) L. Pojman, 310-21, Belmont, CA: Wadsworth.

Wilson, R. A. (2000). The mind beyond itself. InD. Sperber (Ed.), *Metarepresentations: A multidisciplinary perspective* (pp. 31–52). New York: Oxford University Press.

Wilson, R. (2004). *Boundaries of the mind: The individual in the fragile sciences: Cognition*. New York: Cambridge University Press.