

# A representational account of VR immersion

Jonathan Cohen and Matthew Fulkerson \*

There will always be circumstances in which an illusion rendered by a layer of media technology, no matter how refined, will be revealed to be a little clumsy in comparison to unmediated reality (Lanier 2017, p. 49).

---

After a hard day of grading undergraduate exams, you decide to reward yourself with an hour of VR entertainment. You strap on your VR headset (e.g., <https://www.oculus.com>), fire up the undersea exploration game Ocean Rift (<http://www.ocean-rift.com>), and undergo an immersive perceptual experience in which it seems to you that you are exploring and interacting with a green turtle in an underwater environment.

In what follows, we focus on a specific feature of this sort of experience, namely its felt quality of immersion. The feeling of immersion associated with VR experience is arguably a major source of the recent interest in the medium for users and theorists alike. It has seemed plausible to many, for instance, that the immersion afforded by VR distinguishes our experience with this medium from our experience with more familiar depictive representations (e.g., drawings, paintings, photographs, and films). If this is right, then this quality of the medium and our experience of it poses an important *prima facie* challenge to a simple and attractive conception of VR as a representational medium. Though there are many reasons to favor this simple and attractive account — and, indeed, to regard it as our default theoretical description of the medium, it can seem that the view has difficulty accounting for the phenomenon of VR immersion. We will argue that this challenge is superable — that, appearances to the contrary notwithstanding, the representational understanding of VR affords a compelling and theoretically rich account of VR immersion without the need for extravagant additional theoretical resources. If we are right, then the phenomenon of VR immersion is not a reason to give up on the representational view.

---

\*Department of Philosophy, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0119; [cohen@ucsd.edu](mailto:cohen@ucsd.edu), [mfulkerson@ucsd.edu](mailto:mfulkerson@ucsd.edu)

Here is our plan. We'll begin by laying out and motivating a representationalist account of VR and VR experience (§1). We'll expand on this view by attempting to characterize just what content VR might be thought to represent in a range of cases of interest, and whether and in what ways such contents should be counted as falsidical (§2). We'll then begin to consider the phenomenon of VR immersion and the theoretical challenges it raises for a representational account (§3). This will lead us to consider (§4) whether the answer to these challenges is to enlist the thesis, defended by recent authors, that VR representations are *transparent*; we'll argue that the transparency thesis is both undermotivated and unhelpful in accounting for immersion. In response, we'll put forward our own account of VR immersion in terms of fundamental features made available by the simple representationalist framework (§5). Finally, we'll conclude (§6).

## 1 VR representationalism

We begin with what we take to be a platitude about VR that, we suggest, should form the basis for our theoretical description of its nature: like photographs, paintings, movies, and the like, VR is a technology for perceptually representing objects, properties, and events. In the case described at the opening of the paper, both VR and VR experience (viz., the experience a user has in interacting with the VR device) represent a green turtle.<sup>1</sup> This starting idea, simple as it is, turns out to be surprisingly powerful in understanding the properties of the medium of VR, and brings with it several significant explanatory benefits.

### 1.1 Motivating a representational view

The representational understanding of VR explains what might otherwise be puzzling features of the medium and the experiences it supports in terms of familiar features of representation. For example, this analysis affords a simple explanation of the widely shared intuition that the perceptual states we undergo when employing VR are ordinarily *illusory* (Juul 2005; Tavinor 2009; Bateman 2011; Meskin and Robson 2012; Velleman 2013). Namely, the claim that VR represents (aspects of) the world (or at least aspects of the 'virtual world'; more on this shortly) as being particular ways allows us to say that, like other forms of representation, it has content — that it characterizes the (virtual) world, and that such characterizations/content can be assessed for accuracy along different dimensions. Thus, we might think that what the system presents to you on an occasion can be accurate in respect of the color and form of green turtles (green turtles really do have those colors and forms), but inaccurate in respect of the

---

<sup>1</sup>One can imagine both particularist (VR represents a singular content about a particular green turtle) and generalist (VR represents the general/quantificational content that there is some or other green turtle) versions of this thesis. We won't attempt to resolve this matter here. For parallel debates about ordinary perceptual representation unmediated by VR, see Hill (2023) and Logue (2023).

egocentric location of what it represents (there is not a green turtle two meters away and to your left).

Another non-trivial fact about VR easily explained by a representational account is its surprisingly extensive scope. Why is it that if VR can present to its users a red fish, it can also present to its users a blue fish or a red turtle? This point shouldn't be overstated: there are clearly some interesting limitations on what we can experience under VR, presumably due to constraints on perceptual systems, VR technologies in use, and their interactions. Still, and within those limitations, it is neither trivial nor *a priori* that the scope of what can be presented via VR extends as widely as it does. Again, this fact is predictable on a representational account. The scope of possible VR scenarios is as impressively broad as it is because it mirrors (within interesting limits) the scope of possible representational contents: if something is a representable content, it is (in principle) a candidate for being a representable VR content.

But if construing VR in representational terms offers genuine explanatory benefits such as those we've indicated, it is important to see that these benefits come at no new costs (beyond those all of us will have paid long ago). We take it that the posit of representations has more than earned its place in philosophers' toolboxes for the light it sheds in understanding all kinds of features of cognition, perception, language, action, and much more. If, as we are suggesting, informative explanations of important features of VR are made available by appeal to the very same already and abundantly justified posit, we should welcome them.

## 1.2 Representational VR and the extension of perception

We now want to point out a further important attraction of the representational account of VR: it explains how VR can (like other representational media) serve the useful function of presenting individuals and properties when these are not present/instantiated in ways that would allow them to be accessed by unaided perception not making use of VR (cf. [Macpherson 2020](#)). They might be too spatiotemporally remote, or smaller/larger than the spatiotemporal limits of unaided perception, or in environments that would be difficult/dangerous/expensive to access, or might not even exist at all. That VR can be used in these situations is, therefore, often of great practical and epistemic utility.

To see how the representational explanation of this fact about VR goes, consider an ordinary case in which one straps on a VR device. We take it that, in such an ordinary case, one veridically visually represents the display (we'll expand and defend this claim in §1.3). Such veridical visual representation will have as one of its contents the VR display and its features. Crucially, however, representation of the VR display is not the full story of what happens to perceptual systems when engaging with VR. On the contrary, and because the device is itself a representational vehicle with a depictive content of its own — as it might be, content involving a green turtle — in representing the device, one gains indirect perceptual contact with the content of the display. The

content of the display (in this case a green turtle) becomes a second, additional layer of represented perceptual content.

We are warranted in thinking of the connection one secures in this way as perceptual in some good sense; after all, it clearly exploits perceptual (e.g., visual) capacities. Moreover, since the resulting connection affords a link to, and provides content about, something distinct from the first layer of perceptual content, it should be counted as an *extension* of ordinary visual representation.

That said, the connection between the subject and the green turtle in such cases is clearly *indirect*: the subject visually represents the device while the device represents the green turtle, hence the subject-turtle connection runs through an intermediate (the device). This contrasts with the subject-device connection, which, we are assuming, does not causally or metaphysically depend on the subject seeing some other entity intermediate between herself and the device. Now, in saying ordinary perception does not involve perception of intermediary entities, we are rejecting a distinguished representative realist tradition that construes ordinary perception in just such indirect terms. However, we are suggesting that the perceptually mediated contact subjects obtain with the representational contents of VR devices (and those of other representational media) *is* indirect in just this sense — viz., it is indirect in just the sense in which representative realists held that ordinary perception was indirect.<sup>2</sup>

As we have said, such indirect perceptual connections can be extremely useful. Unless your living room is very different from ours, there is not a green turtle swimming in it before you. Hence, the second, indirect, layer of perceptual content (to the effect that there is a green turtle) is not veridical — or if it is, this is a lucky instance of veridical hallucination. But this is just the point of the technology. It allows a perceptual/epistemic connection (indirect and falsidical though it is) with a green turtle when there is no green turtle in the vicinity. The observation that representations are capable of having as contents absent individuals/uninstantiated properties is entirely familiar (cf. [Brentano 1874](#)). In the present setting, that familiar idea about representations sustains the idea, mentioned above, that VR devices are like other representational technologies in allowing for the practically valuable perceptual presentation of individuals and features that would be otherwise perceptually unavailable (or available only at greater effort/inconvenience/expense/danger).

---

<sup>2</sup>Notoriously, such indirect accounts of ordinary perception invite the threat of veil-of-ideas skepticism: if perception directly connects us exclusively with intermediaries numerically distinct from the distal items in the world, then one might think that even successful instances of perception are incapable of ruling out the possibility that we are massively deceived about the character of the distal. Whatever view one takes about the significance of that threat to representative realism about ordinary perception, the analogous skeptical threat with respect to an indirect account of VR strikes us as far less alarming, if only because no one thinks of VR as a fundamental source of epistemic contact, or one that underpins knowledge generally speaking. Indeed, skepticism with respect to the content of VR is appropriate: given that, as we'll discuss in §2, much of this content is, in fact, falsidical (typically by deliberate construction), it is not inadvisable to regard VR content with a healthy dose of skepticism.

### 1.3 Do we veridically represent the VR display?

We said above that users of VR devices visually represent the VR display. We choose this phrasing carefully to sidestep issues about whether a subject has a full-blown perceptual experience of the display (i.e., ‘sees’ it in a robust sense that entails focal attention or capacity to individuate against a surround). We find this a natural view. Moreover, this view easily generalizes to the representation of other features made available by such devices. In a modern flight simulator, for example, a subject will visually represent many screens, but they can also tactually represent various switches and control surfaces, represent a variety of sounds, motion, and even smells (from a simulated engine fire, say). However (philosophy being what it is), we can imagine that some would reject this natural view in favor of the even weaker view that users (/their perceptual systems) are merely causally influenced by the display, but that they don’t count as visually representing it (much less fully *seeing* it). Now, if push came to shove, we could accept this weaker position: the representationalist account of immersion we will offer in §5 could be recast in this form without compromising our main claims. That said, we find the moderate representational view well-motivated, and the apparent reasons for favoring the weaker view unpersuasive.

Why is the moderate view that users perceptually represent the VR interface well-motivated?

For one thing, it is, presumably, much harder to deny that larger and further away representational devices — say, photographs, television screens, computer monitors — are visually represented by their users. Given this, it seems somewhat *ad hoc* to claim that seeing somehow stops when the device falls below some particular threshold of size or distance.<sup>3</sup>

Further support for the claim that we visually represent the display/interface comes from a long line of empirical research demonstrating that subjects represent proximal as well as distal magnitudes when viewing objects — even when, as is typical, those objects are presented on experimental display devices. They visually represent the respect in which the tilted coin looks oval as well as the respect in which it looks round, the respect in which the white wall under red light looks red as well as the respect in which it looks white, the respect in which the telephone pole in the distance looks shorter than a nearer telephone pole as well as the respect in which they look the same size, and on and on (e.g. Arend and Reeves 1986; Cohen 2008, 2010, 2012a,b; Morales, Bax, et al. 2020; Morales and Firestone 2023). Importantly, the finding running through this empirical tradition that subjects can behave (e.g., sort, adapt) on the basis of such proximal magnitudes gives reason for thinking that the latter are not merely causal influences, but are perceptually represented. And, to

---

<sup>3</sup>Here we ignore VR systems that omit displays and project light directly onto users’ retinas (cf. [https://en.wikipedia.org/wiki/Virtual\\_retinal\\_display](https://en.wikipedia.org/wiki/Virtual_retinal_display)). In these cases we allow that there is no perception of the display (because there is no display). We are inclined to think of these cases as hallucinatory representations involving deviant causal chains. Similarly for real or imagined cases involving VR by direct cortical stimulation that bypasses perceptual transducers entirely. We ignore this complication in what follows.

repeat, all of this is true when the stimuli are presented on experimental display devices (e.g., computer monitors). It's hard to see how that could be true if they were merely being causally influenced by those experimental displays. To be fair, we are not aware that such experiments have been replicated using VR displays in particular, but it would be surprising indeed if it turned out that visual representation works in a qualitatively different way when the display gets smaller and closer.

Finally, our view is strengthened by consideration of more modern VR technologies that leverage modalities beyond sight. The idea that a subject in a VR scenario isn't actually representing the sounds, representing the pressure and motion, or representing the odors of the interface, but is instead merely causally influenced by these, is, we think, extremely implausible. When the yoke in a flight simulator begins shaking violently (simulating a stall warning), it certainly seems that the subject represents the interface via touch as shaking. Indeed, we think even the stronger view that they fully perceive the interface is quite plausible here and in many other cases. For many VR systems the subject seems to perceive the device, as fully as if the device were real. Our position does not require this stronger view, but its plausibility in a great many cases lends additional support to the moderate representational claims made here.<sup>4</sup>

All of this is to say that our initial position, to the effect that users of VR devices perceptually represent the VR interface, seems overwhelmingly plausible. On the other hand, as we say, we remain unmoved by considerations that might seem to motivate the weaker alternative view that users are causally influenced by but do not perceptually represent the display.

A consideration against our stronger view that we perceptually represent the VR interface (suggested to us by Robert Briscoe) comes from the observation that the standard range of cues from which the visual system computes the depth of perceived objects is not available in cases of normal use of VR displays (for discussion, cf. Briscoe's contribution to this volume). We take this point, but do not see that it is a reason for denying that the VR user visually represents the VR display. To think otherwise, we would need a reason for believing that the exploitation of the visual depth cues in question are necessary conditions for visual representation (rather than cues that are, in fact, available in many ordinary cases).

A second reason for denying that users of VR devices veridically represent the display comes from considerations raised by [Macpherson \(2020\)](#) (cf. [Macpherson and Batty 2016](#)).<sup>5</sup> Of course, all sides allow that one's visual system tracks information from the VR display. Despite this, it is also notable that in normal use one is not aware of the spatial boundaries of the two screens before

---

<sup>4</sup>An advantage of the more moderate view is that it allows us to sidestep worries about full-blown perception like those discussed by [Hill \(n.d.\)](#). While we ultimately believe such worries can be met, it is an advantage of the moderate claim that we can avoid such hassles.

<sup>5</sup>To be clear, here we are applying Macpherson's considerations about the possibility of veridical perception in VR of various entities (e.g., the 2d image on the screen, depicted objects, the properties of those objects) to our own, distinct, question of whether the VR display itself is veridically represented.

one's eyes that comprise the display, and indeed the two are visually fused into one. For Macpherson, such cases are instances of *object illusion* (which she contrasts with property illusion). Here is what she says about the related question of whether one veridically perceives (not the display, but) the image on the display:

What we have here is a case that is the reverse of double vision: it is half vision! For every two objects present on the screen you have an experience of one object. As your experience is sensitive to the images, we have good grounds for holding that you perceive them—but, . . . qua objects of perception, in an illusory fashion (Macpherson 2020, p. 25).

Parity of reasoning would seem to suggest, analogously, that one does not veridically represent the VR display: here, too, one's experience is sensitive to the display, but represents one object when two are present, so represents the display, qua object of perception, in an illusory fashion.<sup>6</sup>

Once again we are unpersuaded. For one, simply closing one eye while wearing VR goggles does not undermine the virtual experience, and it seems odd to maintain that the overall visual representation goes from veridical to illusory and back as we open and close one eye. More importantly, the consideration about optical fusion only applies to goggles with separate screens. Many VR systems do not have this feature (most flight simulators, for instance, involve large unified screens at some distance from the eyes). If we restrict our claims to dual screen VR devices, holding that such representations of the devices are illusory would seemingly commit us to saying similar things about watching movies through 3D glasses or looking at any autostereograms (i.e., MagicEye images). The moderate representational view we advocate offers straightforward pronouncements about such cases, and in ways that line up nicely with the standard accounts of more traditional representational devices such as photographs, television screens, and computer monitors (viz., we come to represent the depicted scenes in such media by veridically representing the medium itself). Moreover, the same view extends straightforwardly to cases in which the representational devices are so large or so close to our receptors that we lose the ability to represent their edges or source locations in perceptual experience (e.g., visual cases involving large format film screens such as IMAX, or auditory cases involving complex speaker arrays using surround sound technologies, such as those used in movie theaters).

In this section we have put forward a construal of VR as a representational medium, shown that this construal has some non-trivial explanatory benefits, and that it offers the hope of understanding the ways in which VR can extend our perceptual powers. Namely, it offers the hope of seeing VR as supplying

---

<sup>6</sup>One question this raises is why Macpherson thinks misrepresenting the number of items (as opposed to the color, form, or other features of items) breaks the perceptual representation relation. We put this question to one side.

an indirect, though perceptually mediated, form of access to entities and events that would not otherwise be perceptually accessible — either at all, or, else not without greater cost, time, or danger. In §2 we will elaborate on the representational construal of VR, showing that it makes available plausible and theoretically interesting descriptions of a range of cases in which the technology is typically used.

## 2 What VR represents: Some case studies

We have said that VR is a representational medium — that users veridically represent a device’s interface, and that, because the interface is a representational vehicle with content of its own, they are additionally caused to perceptually represent that embedded content. We have said that the first of these layers of content is veridical, while the second, in general, is not. Now, if this is so, one might wonder: what is the second layer of representational content in VR experience? In this section we want to present a cluster of answers to this question with respect to canonical types of VR experience. We do so not in order to select and defend any particular favorite answer, but rather to show the wide and varied landscape of plausible options available to representationalists. Characterizing the representational content in salient uses of VR will also put us in a position to say whether and in what respects this second layer content is falsidical in such cases.<sup>7</sup>

### 2.1 VR Simulation

Many VR experiences aim for a kind of realism. One of the most salient approaches involves mixing virtual elements and actual reality. These cases go by the label *augmented reality* (*/AR*); they can produce virtual experiences that are embedded or connected to actual objects, spaces, and properties in our perceptual environment. AR experiences can consist of digital overlays which appear on top of actual objects or they can add virtual objects to the local environment through various displays and interfaces. In these sorts of scenarios, we can say that the VR device presents objects and properties in the actual environment that are not actually there. While there are genuine entities in these mixed cases, and these elements can be veridically perceived, it is the presentation of non-existent objects and properties that generates illusion. When playing *Pokemon Go*, for instance, we might experience a Charizard on the screen, apparently running around in our actual backyard, which we otherwise perceive as normal. While the (direct and indirect) awareness of the actual yard may be accurate, we do not need to say the same thing about the Charizard. On the representational view, there is no actual object in our

---

<sup>7</sup>For simplicity, in this section we work with a traditional framework for thinking about illusion and hallucination in representation, rather than the more fine-grained characterization offered by [Macpherson and Batty \(2016\)](#) and applied to VR by [Macpherson \(2020\)](#). None of our substantial theoretical commitments below depend on this choice.

local environment that corresponds to the Charizard representation. It is simply erroneous. The representationalist is free to treat these elements of AR cases as garden variety illusions/hallucinations. In these cases, there are real world objects and spatial locations that are being assigned features they don't have (illusions) or are experienced as containing objects that do not exist (hallucinations), resulting from the use of an overlay screen or other interface.

Some VR experiences aspire to realism without incorporating any real objects or properties. Call these examples of *VR Simulation* (to be contrasted shortly with *VR Fantasy*). Thus, in the Ocean Rift case with which we opened, the VR visually represents a green turtle in the space around the user when there is, in fact, no green turtle. Such systems might nevertheless generate experiences of real objects, places, or properties.

Consider a modern flight simulator, for example, which might include dozens of screens, lights, and controls, all of which individually and collectively match what we would see in an actual aircraft. The representationalist will treat each of these elements in the standard way: as representations of actual objects and properties. On this view, the interface merely represents a real world object. The fuel gauge in a flight simulator is not a real, functioning fuel gauge, it just looks and functions like one. When it shows that the gas tanks are half full, this is a representation of a non-actual state (there are no gas tanks in the simulator).

On this conception, VR simulations can be illusory in representing the world in ways that it is not. This applies even when the apparatus is intended to represent a real world feature or individual. The simulated fuel gauge is designed to mimic one in the real world, and its functions and content are meant to convey features in the actual world. But neither actually does so; both the gauge and its content are illusory, representations of items which are not actual. This point generalizes to each object and property experienced in a VR simulation.

## 2.2 VR Fantasy

Not every VR system attempts to replicate or simulate actual objects or locations. Some involve entities that do not actually exist in the physical world. We will call these cases of *VR Fantasy*. Note that the boundaries between these classes are fluid. Many VR experiences will involve a mix of simulated and fantasy elements, for instance.

Let's consider an example. In the VR fantasy world presented in Spielberg's film adaptation of Ernest Cline's *Ready Player One*, there are virtual stores that run on virtual currencies, and virtual objects that can be tracked, stored, sold, and manipulated. These elements seem to have a persistent multi-subjective reality.<sup>8</sup> These features can initially make it seem more difficult to say what

---

<sup>8</sup>Some of these might resemble actual stores or currencies. Indeed, the virtual objects might interact with objects outside the fantasy world. For example, one can spend real world currency on the virtual objects available in Second Life (<https://secondlife.com>). Such interactions are not a problem for the representational view. It would be no different than spending real money on any other items for sale in representational media. For instance, someone playing a tabletop

counts as the content — and thus what would count as veridical or illusory — in such cases.

However, there are a number of plausible accounts available to the representationalist. One possibility is to assimilate Fantasy cases to instances of VR Simulation. On this approach, the VR medium falsidically represents the virtual store and its virtual contents as bearing real-world features and locations with respect to the user. Understood this way, VR fantasy poses no special difficulty for the representational account: the account would treat such cases as instances of illusory/hallucinatory representation by the medium in just the way it treats the entities experienced in a flight training simulator as illusory/hallucinatory.

A second possibility is to treat the contents of these experiences as only involving *virtual* objects and properties, each located not in real locations but virtual ones. On this view, the represented content might indeed turn out to be veridical. When a VR headset is producing an experience of a large mountain in the distance, the VR system represents, of the *virtual* space presented in the system, that there is a *virtual* mountain in the distance. This content would seem more plausibly true than the parallel content about a real, actual mountain in the distance. Even so, the representationalist view would not be undermined by such a move. They could allow that, given our understanding of virtual objects, it is indeed true that a virtual mountain is being represented in the distance, just as a carved wooden figurine can veridically represent a queen when placed on a chessboard. This view is not required, however. The representationalist could also maintain that the virtual content is false in a straightforward way: virtual properties and virtual spaces are not real. It would follow that there are no virtual mountains or virtual spaces around the subject. There are properties and spatial locations being *represented* by the media, but these are not real entities and representing them as such does not require we treat these expressions as veridical. Given this point, it would be reasonable for the representationalist to treat the content of VR Fantasy as erroneous, but more in the manner of a hallucination rather than an illusion.

There is a third, final option, which would be handled in the same way. Instead of introducing a novel category of virtual objects and properties, we could simply treat Fantasy VR entities just as we treat fictional entities in other representational media. While there is no actual Sherlock Holmes or 221b Baker Street, they do exist in Conan Doyle's fiction. Similarly, there are no (actual) virtual mountains, but they do exist in the fictional worlds presented via VR (cf. [Chalmers 2017](#); [McDonnell and Wildman 2019](#)).

As above, a representationalist tempted by this line can maintain that VR is indeed veridically representing the existence of fictional objects. Or they could reject this reading, and maintain that, strictly speaking, such contents are false (since fictional entities don't exist). There are intense ongoing debates about how to understand the ontology of fictional objects and the semantics of expressions that embed fictional operators ([Lewis 1978](#); [Thomasson 1998](#);

---

version of Dungeons and Dragons might offer real money for a new custom figurine, and someone may lose real money betting on the outcome of a video game tournament.

[Kroon and Voltolini 2023](#)). But, as above, the representational view itself does not force us to choose among these options. In treating VR experiences as yet another representational media, no different in kind from films, paintings, and traditional video games, the representationalist will treat fictional VR elements like any of the others.

As we've seen, VR experiences can go out of their way to accurately represent real world objects and locations. In professional flight simulators and VR surgery training devices, every care is taken to make the experience match as closely as possible what the corresponding actual perceptual experiences would be like. As technology improves, we can imagine the noticeable differences becoming insignificant. Even in these cases, we believe that the representationalist has a clear and plausible account of the VR contents: they are illusory/hallucinatory. Other VR experiences make no attempt to replicate real world experiences. In these systems, entities and locations are entirely made up. In such cases, the representationalist is free to treat VR Fantasy as any other fictional medium.

### 3 Immersion: A problem for representationalism?

We distinguish between *VR immersion*, the interesting and putatively distinctive quality of the experiences we undergo when interacting with VR, and *VR immersiveness*, the capacity of VR, qua medium, for inducing immersion. Both deserve explanation.

We begin with VR immersion, which is perhaps the most significant and oft-remarked-upon feature that distinguishes VR experience from our experience of other representational types, and a large part of the interest and attraction of the medium. It is difficult to characterize this sense of immersion in precise terms.<sup>9</sup> Still, we take it that this feature of VR experience is salient and familiar even to casual users. A subject embedded in a well-constructed VR environment feels as though she is embedded in a genuine space, with various objects located around her. Even very basic systems without haptic feedback (say, systems involving cheap cardboard goggles holding a cell phone) can generate a strong feeling of immersion. The most sophisticated systems (e.g., those involved in Level D flight simulators used to train airline pilots, which typically involve accurate cockpits and controls, high resolution screens, six degrees of freedom, along with simulations of actual airport traffic, weather, ATC communication, and various aircraft technical scenarios) can generate an even stronger sense of immersion.<sup>10</sup>

---

<sup>9</sup>It has been variously understood in the literature in terms of a "feeling of presence" ([Matthen 2005](#)), the "perceptual presence" that arises from the felt preservation of sensorimotor contingencies ([Noë 2004](#)), or the immersiveness that arises through processes of sensory interaction ([Aurelia 2025](#)). While these are all important and interesting features of VR experience, in our view, all of them flow from something more fundamental, and that goes to the heart of VR's representational character.

<sup>10</sup>The second season of the television show *The Rehearsal* includes a particularly nice examination of flight simulation (among other forms of simulated reality).

On its face, the phenomenon of VR immersion/immersiveness, might seem to pose an obstacle to the sort of representational view of VR we have offered. After all, immersion seems to be absent or else present to a significantly lesser extent for other paradigm representational media (from comics to movies to photographs), and it is not obvious at first glance why things should be so saliently different for VR if the latter is like all of those other media in being just another form of representation. Indeed, and as noted above, the very structure of the representational view, on which the object of perception is a mere representation (rather than, say, an actual green turtle) might seem to require something less like an immersive connection between subject and object, and more like an intervening veil of perception between subject and object. As we shall see (§4), this thought might reasonably be thought to motivate alternative theoretical approaches designed to allow for a much more direct, perceptual/epistemic relation between subject and object. We believe this is a mistake, and that the representational view we have proposed is, perhaps contrary to initial appearances, well-suited to account for the phenomenon of immersion all by itself, i.e., without the need for further, and more controversial, theoretical tools.

Before we turn to proposed explanations of immersion/immersiveness, we want to register a historical caution against treating the latter as a reason to forgo explanations in terms of representations. While VR immersiveness has led some to think of VR as beyond the reach of representational explanation, a casual review of the history reveals that consumers were similarly impressed by the immersive capacities of many other newly introduced media that have, with the passage of time, come to seem much less controversially representational. Indeed, nearly every introduction of a new representational technology brings in its wake experiences that seem authentic and accurate, embedding their subjects in new forms of awareness that are often described as realistic. The introduction of linear perspective allowed for paintings that were more immersive and ‘realistic’ than what had come before. In a recent article highlighting the invention of photography, [Sullivan \(2018\)](#) tells the story of a young boy who saw a horse and carriage projected on his ceiling when the small gap in his curtains — along with the perfect outdoor light conditions — accidentally turned his room into a *camera obscura*. His parents initially thought he was going mad, with his talk of seeing horses and people on his ceiling. But the experience was vivid and immersive: “In the seconds when the pictures had danced above his head, it had seemed that, had he been able to reach, he could hold them in his hand, like a print” ([Sullivan 2018](#)). From the camera obscura we move to actual photographs, and from there even more immersive technologies. In the (relatively) rapid transition from still photography that occurred between the introduction of the Daguerreotype to the introduction of full-fledged film or ‘moving pictures,’ there were a number of intermediate and highly immersive new technologies introduced, including Edison’s Kinetoscope, the Bioscop, and the Lumiere’s Cinématographe (cf. [Aurelia 2025](#), ch. 1). Audiences at each are often described as mesmerized

by the seeming reality and immersiveness of the technology. Most famously, of course, audiences are said to have reacted in shock and fear, screaming and running to the back of the room, during the first showing of the Lumiere's *L'arrivée d'un train en gare de La Ciotat*, wherein a train is shown heading straight for the audience.<sup>11</sup> Similar claims of immersion/immersiveness can be found for the later introduction of sound to previously silent films, the introduction of color and 3D images, the use of computer generated images (CGI), and the use of large format screens like IMAX. Even this briefest of histories of representational technologies encourages caution about attributing to VR a unique ontological or experiential status simply on the grounds that its early adopters find it especially immersive or realistic.

We take this historical cautionary tale to heart. Still, the frequency and force with which commentators point to the special immersive quality of VR experience suggests that there is something here worth explaining.

## 4 Transparency and immersion

Some theorists have held that VR devices are “transparent” (/that they are “visual prosthetics”) — i.e., that they not only afford an indirect and perceptually mediated informational connection between user and represented object, as we are claiming, but that their use counts as a way of perceiving/seeing the represented object. This view is typically presented as an extension of Walton's deliberately heterodox view that photographs are visually transparent — that when we see a photograph of Granny, we thereby additionally stand in the seeing relation to (i.e., we see) Granny (Walton 1984). Whatever one thinks of the plausibility of this view, it is worth taking VR transparency seriously in the present context because, whatever other attractions that view may hold, one might think it offers the prospect for answering our puzzle about VR immersion/immersiveness — that the distinctively immersive quality of our VR experience of a green turtle could be explained partly in terms of the transparency of the medium. Moreover, precisely because the envisioned explanation would cite a feature of VR that goes beyond its representational status, one would additionally be in a position to say why VR's immersiveness is not shared (or not shared to the same extent) by other representational media.<sup>12</sup>

---

<sup>11</sup>This well-known report has been subject to scholarly controversy. [Loiperdinger and Elzer \(2004\)](#) doubt its truth altogether. Wikipedia credits to Benjamin Bratton (but without citation) an alternative explanation for the audience's reported reaction — viz., that audience members thought the images were being produced by a camera obscura, which *would* have been strong evidence for an actual train barreling down on them.

<sup>12</sup>Though we intend the main discussion of this chapter to be officially neutral on the question of VR transparency, we register for the record that we ourselves reject the claim that photographs (and, by extension, VR devices) are transparent. Briefly, we are sympathetic with the suggestion ([Cohen and Meskin 2004](#); [Meskin and Cohen 2008](#)) that standing in a perceptual relation to an object *O* requires not just a causal/reliable tracking relation, but the carrying of information about

Granting VR transparency *arguendo*, how would this help in explaining VR’s immersiveness? The idea would be that VR devices afford immersion because, in using them one is thereby doing more than merely receiving information about the green turtle — one is literally, directly *perceiving* (/seeing) it. And, while there is, of course, disagreement about the details, many theorists would accept that perception has some sort of distinctive epistemic and phenomenological status lacked by other informational channels. Perhaps, then, this distinctive epistemic and phenomenological status, rooted in genuine perception rather than mere information carrying, is enough to make good on the immersiveness of VR.<sup>13</sup>

We want to argue that this proposed account of VR immersiveness is unsuccessful. We’ll do this by considering two flavors of the transparency thesis for VR — a first on which the represented objects that transparency claims we perceive is an ordinary if spatiotemporally remote concreta such as a green turtle (Dilworth 2010; and, at least for VR tokens involving live-tracking, Tavinor 2019), and a second on which the represented objects that transparency claims we perceive is, instead, a “virtual object,” such as a digital avatar that turns out to be identical to a data structure in a computer (Chalmers 2017, 2022, ch. 10–11; cf. Chalmers 2005). We’ll label the first flavor of the view *ordinary object transparency*, and the second flavor *virtual object transparency*.<sup>14</sup>

A first, internal, problem for the proposal for explaining immersiveness in terms of either flavor of transparency depends on how much non-VR

---

the egocentric spatial relation between perceiver and  $O$  — a requirement satisfied in unmediated instances of perception and perception via genuinely transparent media such as mirrors and telescopes, but not when one perceives a photographic or VR representation of  $O$ .

Macpherson (2020) offers a number of objections to this further requirement, one of which is that she finds it *ad hoc*. This leads her to hold that “there is no privileged information that must be transmitted by a process of seeing or an experience of seeing in order for that experience to count as one of seeing” (p. 23). Indeed, in conversation she has suggested that there is *no* principled requirement on the perceptual relation beyond the reliable tracking by the perceptual state of  $O$  and its properties (cf. Macpherson 2014). While (tragically!) reasons of space prevent us from defending Cohen’s and Meskin’s proposal from her objections here, we do want to suggest that the no-further-constraints view (which of course can be used to motivate VR transparency) is too liberal. For one thing, the state of the visual system arguably reliably tracks many intermediate stages on the causal pathway from  $O$  to the perceiver, so the unconstrained view allows that we perceive all of these, and all of them should additionally count as transparent perceptual prosthetics. For another, unless one allows that all of these intermediate states are part of the content of perceptual experience (which seems deeply implausible), one needs some condition by which to distinguish those reliably tracked objects that are part of the content of perceptual experience from those that are not. But it’s unclear why a condition of that sort should be acceptable if (as proposed) there is no principled condition that distinguishes reliable tracking that counts as perceiving from reliable tracking that does not. These considerations suggest to us that an adequate view of object perception will require at least *some* further constraint beyond mere reliable tracking (of which Cohen’s and Meskin’s proposal is but one candidate).

<sup>13</sup>To be clear, the use of transparency to explain immersiveness we explore here is conjectural: the VR transparency proponents discussed below do not endorse this (or any other) account of immersiveness as such.

<sup>14</sup>We note that Macpherson (2020) does not count as a proponent of ordinary object transparency. She holds that we see both the distal turtle and *features* of the display; but, crucially, she denies that we see the display itself. Hence it’s not true, for her, that we see the ordinary object by seeing the display.

transparency one is willing to accept. Recall that our explanatory target is an allegedly distinctive feature of VR, as contrasted with other representational media, such as photographs, television, and realistic paintings. Theorists will differ over just which of these non-VR examples of representational media count as transparent; but insofar as the motivation for VR transparency stems from extending to VR a view about photographic transparency, it is unlikely that transparency will turn out to hold distinctively of VR. But if so, then the appeal to transparency is unlikely to provide an explanation of why immersiveness turns out to hold distinctively of VR. Indeed, as we've seen, advances in representational technologies have often been credited with an increase in immersive quality. Appeal to transparency in the case of VR would seemingly force us to either adopt a distinct account of immersiveness in these other media or extend transparency to video games, movies, and other pictures. Either option is unpleasant. If we accept a non-transparency account of immersiveness for movies and video games, it's unclear why that account should not be extended to a non-transparency account of VR immersiveness. And while transparency might initially seem appealing for VR, it is, we think, much more problematic for other representational types.

A second problem with the proposal to understand VR immersiveness in terms of transparency arises for the ordinary object flavor of the view in particular. The problem here is that the proposed analysis has a narrower scope than the analysandum. We can make this point by observing that, as noted in §1, VR representation (like other representation) extends not only to "good cases" where the medium represents objects that exist (say, a green turtle) but also to "bad cases" where the medium represents objects that do not exist (say, Pegasus/a pegasus in Chicago), and that there is no salient difference between the two in terms of immersiveness. In, say, a fantasy VR case, our (good case) VR experiences as of existent objects can be exactly as immersive as our (bad case) experience as of non-existent objects. However, the thesis of ordinary object VR transparency is at best applicable only to good cases. The ordinary object transparency theorist will hold in the good case that use of the VR device affords a genuinely perceptual connection to a green turtle. Similarly, ordinary object transparency in the bad case would require that by perceiving the device, one perceives Pegasus/a pegasus in Chicago — i.e., standing in a genuinely perceptual relation to Pegasus/a pegasus in Chicago. But that's not possible. One can't stand in a genuinely perceptual relation or any other relation to Pegasus/a pegasus in Chicago, because (at the risk of disappointing readers) Pegasus/a pegasus in Chicago lacks the perfection of existence. Since the phenomenon of immersiveness extends to cases where ordinary object transparency fails, there is no hope of explaining the former in terms of the latter.

The other flavor of VR transparency, virtual object transparency, sidesteps the problem just raised, because the virtual objects it holds that we perceive are common factors present in both good and bad cases. That said, the attempt to explain VR immersiveness by appeal to virtual object transparency comes with a number of its own problems. First, at least as Chalmers unpacks it, virtual

object transparency rests on a number of controversial commitments, including his realism about virtual objects, his claim that virtual objects (as opposed to the vehicles accepted by any representational account of VR) are themselves perceived, and on accepting his novel, functionalist theory of the perception of virtual objects. Second, and even worse in the present context, adopting these elements as a way of providing an account of VR immersiveness would be question-begging, given that Chalmers argues for his views partly on the basis of VR immersiveness. In other words, appealing to the rich metaphysical commitments to explain immersiveness gets the order of explanation wrong, since immersiveness itself provides some of the strongest motivation for the metaphysical view.<sup>15</sup> Third, at a more fundamental level, virtual object transparency seems an unpromising route to securing a plausible account of immersion: on its face, it appears that one could stand in Chalmers's favored functional relation to virtual objects whether or not the resulting experience is immersive in the sense we are hoping to explain.

But the most important reason not to appeal to transparency in either its ordinary object or virtual object flavor as a way of understanding immersiveness is that, in this context, it is explanatorily otiose. After all, transparency advocates (of whatever stripe) are already committed to treating VR as a representational medium. Transparency is the view that there is a representation, and that in perceiving the representation one thereby counts as perceiving the numerically distinct content of that representation. Obviously, therefore, one cannot endorse that view (in either form we have considered) without first accepting a representational account of VR. But, as we will argue in §5, the representational account already contains the resources to explain VR immersion and immersiveness all by itself. Thus, if what we say below is correct, then, at least with respect to the goal of explaining immersion/immersiveness, there is no point in adding a further commitment to transparency. As far as explaining immersion goes, the latter is an idle wheel.<sup>16</sup>

---

<sup>15</sup>Transparency is a key plank in Chalmers's case for his iconoclastic claims that, contrary to what he admits are standard views (Chalmers 2017, p. 309, 2022, p. 12), virtual reality is real rather than unreal, and our access to the latter by VR devices is a species of veridical perception rather than illusion. If transparency really does lead to these conclusions as Chalmers argues, we find ourselves tempted to run *modus tollens* where he runs *modus ponens*.

That said, we can't help noting that Chalmers's case for VR transparency rests on his desire to extend Walton's thesis of photographic transparency (a view he accepts more or less uncritically), and on his claim that, in VR, we perceive (veridically) the causal antecedents of our perceptual experiences (Chalmers 2017, p. 318). Both of these commitments are highly unobvious.

For further criticism of Chalmers's view that virtual/digital objects are contents of perception in VR, see Hill (n.d.).

<sup>16</sup>Similarly, a commitment to transparency adds no explanatory value over the strictly weaker representational view in accounting for the practically/epistemically valuable extension of perception by VR discussed in §1.2.

## 5 A representational account of VR immersion and immersiveness

If transparency is not the key to understanding VR immersion/immersiveness, then what is? We claim that a representational account contains, by itself, all the resources needed to understand both VR immersion and VR immersiveness.<sup>17</sup>

### 5.1 VR immersion as use-mention error

We propose that a representational account (with or without the assumption of transparency) allows us to make sense of immersion in terms of a specific type of use-mention error that VR is deliberately constructed to support.

To see how this story goes, begin by comparing two episodes: in the first, there is a veridical perceptual experience as of a green sea turtle; in the second there is a VR experience of a green sea turtle. Both experiences represent a green sea turtle, though in only one of them (the first) is there a green sea turtle present in the immediate environment.

Now, we know, at least in outline, the causal story of how the perceptual representation of a green sea turtle is brought about in the first case: this is just the causal story from perceptual psychology that begins with the impingement of energy from the green turtle on our perceptual receptors, and ends with the formation of a perceptual representation. In the second episode, however, the causal story leading to the formation of the perceptual representation of a green sea turtle does not begin with a green sea turtle in the immediate environment — for there is none. Instead, the causal story begins with *the display on the VR device that is itself a representation of a green turtle*. That is, in the second episode, it is not (as in ordinary cases) perceiving a green turtle that causally leads to the formation of a perceptual representation of a green turtle. It is perceiving a VR *representation* of a green turtle that causally leads to that outcome.

We have said that in both the veridical perceptual episode and the VR-involving episode the perceptual systems represent a green turtle, and that a green turtle is present in the first but not the second. This means that, in the VR-involving episode, the perceptual system is undergoing an ordinary, first-order perceptual error: it represents as present something that is not present. Distinctively, however, this ordinary perceptual error is also a kind of use-mention error. It is an error that arises when the perceptual system treats a representation of a thing (viz., the VR representation of a green turtle), rather than the thing itself (viz., a green turtle), as satisfying its content that a green turtle is present.<sup>18</sup> In so doing, it responds to mere mention in the way apposite

---

<sup>17</sup>To be fair, we don't aim to capture with our account all of the diverse things every reasonable thinker has meant to include under the heading of VR immersion. Among things left out, our analysis says nothing about the idea that VR experiences are more interesting or engrossing than experiences of other types. (Thanks here to Katalin Farkas.)

<sup>18</sup>Fans of transparency can express the idea by saying that, where in the ordinary perceptual case one perceives *only a green turtle*, in the VR-involving case one perceives (or at least perceptually represents) *both a green turtle and a representation of a green turtle*. On this version, the use-mention

to genuine use. This deliberately engineered use-mention error, we claim, lies at the heart of VR immersion.

## 5.2 VR immersiveness: The construction of use-mention errors

We claim that the medium of VR exhibits immersiveness because its instances are deliberately constructed to bring about the sort of use-mention error we claim constitutes immersion. This is no mean feat. It is a substantial engineering task to build a representation of  $X$  that causes perceptual systems to go into a state representing the presence of  $X$  when  $X$  is not present. For example, one of the outcomes one needs to block is the ordinary, non-use-mention-confusing response, i.e., that perceptual systems would respond to a representation of  $X$  by forming a perceptual representation of (not the presence of  $X$ , but) the presence of a representation of  $X$  — which, in fact, is its cause. That metarepresentational perceptual state would be veridical. But the point is that well-constructed VR systems are engineered so as to block that veridical metarepresentational outcome and instead to generate the use-mention-error-involving perceptual representation of the presence of  $X$ . To get this to happen, the general idea is to build VR representations of  $X$  that present data to perceptual systems as similar as possible to the data that would be presented to perceptual systems by  $X$  itself, and so to bring about the perceptual representational outcome normally caused by  $X$ .

Designers/builders of VR systems employ a variety of tools to bring about this outcome (cf. [Simeone 2018](#); [Aurelia 2025](#)). Unsurprisingly, since these tools are designed to exploit properties of our perceptual endowments, they are deeply informed by the science of perceptual psychology. While we can't offer an exhaustive list, it is worth considering some of the most important among these tools.

One obvious strategy is to build devices that obscure traces of their own artifactual status. Thus, for example, visual displays should be presented in ways that hide their edges — so that the visual angles they subtend are at least as large as the visual field. That this strategy is commonly pursued explains why VR visual displays are typically presented at relatively close viewing distance, or conversely, through large projections like those which entirely surround the artificial cockpits in high-end flight simulators.

Another cluster of strategies involves building VR devices that present the full surrounding space around the subject, and allowing for extended, dynamic traversal of that space, as well as interactive engagement with objects within it. The space shuttle simulator at NASA's Kennedy Space Center employs a system like this to recreate the experience of a shuttle launch. As we have seen, early reactions to photographs and motion pictures suggest that something similar applies to these media as well. The strategy is taken further in typical VR devices, however. Spaces in a virtual environment feel extended and

---

error consists in misconstruing an episode in which one perceives/represents both of these two things for an episode in which one perceives only the first.

traversable, and many features of object perception in genuine space such as occlusion, permanence, exclusion, and parallax seem to apply to the things represented by VR. The net effect is that VR representations are designed to provide sensory information of represented targets that closely matches what one would experience in veridical unmediated perception of the target (cf. Aurelia (2025) on “sensory immersion”).

Moreover, VR representations, unlike static images or passively watched films or video, provide the user with opportunities to act (under voluntary control) in ways that are represented perceptually as their moving around objects and interacting with them, all while maintaining a realistic set of relations between our perspective and the represented world. Aurelia (2025) notes that such interaction, and the ability to seemingly causally influence the objects around us in the virtual world, facilitates a feeling of “incorporation” into the virtual world. Thus, in a flight trainer, if the user pushes the stick left, the runway ahead of them shifts in an appropriate manner consistent with their movements. A large cloud in front of the plane occludes objects beyond in a manner consistent with their apparent locations in space, and this changes appropriately depending on the evolving representations of the spatial relations between them and the subject over time. Thus, VR devices are built to mimic the object perception supporting sensorimotor contingencies (e.g., occlusion, permanence, exclusion, parallax, and the like) that are part of ordinary perceptual experience as one moves and acts in space.

When these conditions hold, the VR world and its objects feel to a subject to provide affordances for active, exploratory engagement in a way different from those provided by static images or even moving images on screens. Crucially, VR systems present these affordances by ensuring that the visual stimuli presented actively engage motor systems, providing information about the possibilities for direct exploratory engagement with the virtual world and its objects. The visual information provided is richer than what is provided by static images.

And, of course, though we have stressed visual examples, most VR interfaces also involve the coordinated involvement of other sensory modalities. Auditory information, for instance, is an important element of most VR systems. Some even include haptic and orientation information provided by controllers, gloves, or other sensors located on the body. Flight trainers can involve sounds (from outside the cockpit, and from all the systems within), and bodily feedback through the seats and pedals, vestibular activation through motorized movements of the simulator itself, and even smells and thermal changes (simulating a cockpit fire, for instance). Once we include the coordinated inputs made available through these other modalities, we can expand and enrich our explanation of VR immersion by appeal to the way that coordinated sensory inputs are fused into coherent multimodal representations of the distal environment.

Just to consider one example, focus on how this would work in an audio-visual case. When we are able to visually track a virtual object that appears some distance in front of us, and we are provided with reinforcing auditory

information about that object, the joint operation of those modalities provides an additional layer of immersion. Of course, multisensory interactions involve more than the mere co-activation of distinct sensory networks. Instead, the senses interact at multiple levels to provide richer and more detailed information about object locations, features, and identities. We know for instance that visual information about lip movements can increase our accuracy and resolution for understanding speech in noisy environments ([McGurk and MacDonald 1976](#); [Skipper et al. 2007](#)), that sounds can aid tracking of objects through visual occlusions ([Kvasova et al. 2019](#)), and that in general the joint operation of multiple modalities can increase our perceptual accuracy and engagement. The coordinated interaction among sensory inputs results in super-additive facilitation effects that allow novel forms of perceptual awareness possible only in multi-sensory contexts (cf. [Fulkerson 2014](#); [O’Callaghan 2019](#)). VR systems using more than one sensory modality in coordination can tap into these more immersive forms of sensory information.

Needless to say, the kind of coordination required in such a case is not simply a matter of presenting more distal information about the virtual objects and their apparent locations. These forms of multisensory interaction are deeply connected with felt changes in our awareness of ourselves as situated in and moving through a particular location. The coordinated joint operation of multiple senses plays a crucial role, for instance, in both bodily self-awareness (our awareness of our own spatial boundaries in an environment) and in the construction of peripersonal space (our spatial awareness of the area immediately around our bodies in which we can directly interact) ([Fulkerson 2021](#)). These forms of self-awareness are further enhanced by our capacity to directly engage and interact with our virtual environments, building up an even stronger sense of incorporation.

In our view, then, VR immersiveness is the result of careful engineering. A well-constructed VR that engages these tactics (among others) will obscure as much as possible mismatches between the experience of the VR representation of a target and the ordinary perceptual experience of the target, and thereby will encourage the perceptual use-mention error that makes for VR immersion. Without these mismatches to serve as experiential cues that would pull the two apart, we/perceptual systems are apt to treat experiences involving perceptual mention (as it might be, experiences of a VR representation of a green turtle) in ways appropriate for experiences involving perceptual use (as it might be, experiences of a green turtle).

## 6 Conclusion

We have argued that, though VR immersion/immersiveness might seem to preclude a representational account of the medium, they can be straightforwardly accounted for within a representational view. VR immersion amounts to a particular type of use-mention error arising from the ways in which

VR's purpose-specific, carefully engineered, representational vehicles interact with our perceptual mechanisms. And VR immersiveness amounts to the conduciveness of VR systems to that sort of use-mention-based immersion. Beyond its capacity to vindicate the representational framework for thinking about VR, our account of immersion and immersiveness makes a number of interesting and plausible further predictions.

A first is that, given our story, the immersiveness of vehicles (/vehicle-types) is clearly gradient, simply because representational vehicles (/vehicle-types) can be better or worse in providing to perceptual systems data that matches the data they would receive in veridical unmediated perception. A system that offers more intermodal coordination, or one that respects a larger number of sensorimotor contingencies, is predictably more conducive to the use-mention error we have described, and so, on our account, exhibits greater immersiveness than a system that exploits these tools less fully.

Our story allows a related prediction with respect to the history of immersion we noted in §3. We observed there that commentators on new representational technologies stretching through linear perspective, daguerreotypes, the camera obscura, motion pictures, and more, have frequently pointed to the impressive levels of immersion supported by the new tools (often using the same language found in contemporary responses to VR). This is unsurprising given what we have said. If immersiveness is indeed gradient, then one would expect that new representational cases will be counted as noteworthyly immersive when they exceed levels of immersiveness supported by whatever other representational technologies are prevalent. (Analogy: Fred warrants application of the gradable adjective 'tall' if he falls higher on the associated scale — in this case, height — than the contextually provided threshold for that scale (Klein 1980; Kennedy 2007).) We should, therefore, expect that the advancement of representational technology will gradually increase standards for representational immersiveness, with the result that what seemed notably immersive in an earlier time can seem less so at a later time.<sup>19</sup>

Turning from diachronic to synchronic comparison, our account also predicts differences in immersiveness corresponding to distinct representational types. Thus, for example, we should expect that novels are in general less immersive in our sense (though they can be immersive in other senses!) than most movies: the gap between the data presented to our perceptual systems by the novel's representational vehicles (written words) and those presented by ordinary perception is typically larger than the gap between data presented by the movie's representational vehicles (visual images and coordinated sound) and those of ordinary perception. Similarly, we should expect that (holding other things equal) large-scale paintings are likely to be more immersive than small-scale paintings: the vehicles of the former subtend larger viewing angles than those of the latter, and so better obscure one kind of visual cue that distinguishes experiencing them from experiencing what they depict — namely,

---

<sup>19</sup>Of course, we allow that such evolution could slow or stop at some point, or that a technology could exceed some psychological threshold such that it never ceases to strike users as immersive despite later developments.

the presence of the canvas edge. While there is much more to say with respect to such comparisons of immersiveness between representational types, we suggest that the predictions we have considered are indeed plausible, and that this lends support to the account of immersion from which they emerge.

VR technology is becoming more pervasive and connected to our everyday lives. Many of us spend considerable time each day in virtual spaces, in our entertainment, meetings, and work spaces. These activities increasingly include connections via immersive VR devices that can seem fundamentally different from more mundane representational experiences like movies, television, radio, and photographs. In this chapter, we have argued instead that VR devices are in fact best understood as a representational medium, and that we should think of VR experiences as akin to those generated by photographs, paintings, and films. Despite initial appearances, we have shown that the standard representationalist account offers satisfying accounts of VR immersion and immersiveness, and of the fundamentally illusory nature of the medium.<sup>20</sup>

## References

- Arend, Lawrence and Adam Reeves (1986), "Simultaneous Color Constancy," *Journal of the Optical Society of America A*, 3, 10, pp. 1743-1751.
- Aurelia, Sagaya (2025), *Immersive Technologies: Navigating the Impacts, Challenges, and Opportunities*, Routledge, New York, ISBN: 9781040266670.
- Bateman, Chris (2011), *Imaginary Games*, Zero Books.
- Brentano, Franz (1874), *Psychology From an Empirical Standpoint*, Routledge and Kegan Paul, London.
- Chalmers, David J. (2005), "The Matrix as Metaphysics," in *Philosophers Explore the Matrix*, ed. by Christopher Grau, Oxford University Press.
- (2017), "The Virtual and the Real," *Disputatio*, 9, 46, pp. 309-352.
- (2022), *Reality+: Virtual Worlds and the Problems of Philosophy*, W. W. Norton & Company, New York.
- Cohen, Jonathan (2008), "Color Constancy as Counterfactual," *Australasian Journal of Philosophy*, 86, 1, pp. 61-92.
- (2010), "Perception and Computation," *Philosophical Issues*, 20, 1, pp. 96-124.
- (2012a), "Computation and the Ambiguity of Perception," in *Visual Experience: Sensation, Cognition and Constancy*, ed. by Gary Hatfield and Sarah Allred, Oxford, New York, pp. 160-176.
- (2012b), "Perceptual Constancy," in *Oxford Handbook of Philosophy of Perception*, ed. by Mohan Matthen, Oxford University Press, Oxford.
- Cohen, Jonathan and Aaron Meskin (2004), "On The Epistemic Value of Photographs," *Journal of Aesthetics and Art Criticism*, 62, 2, pp. 197-210.

---

<sup>20</sup>We are grateful for feedback from the editors of this volume, and from audiences at the Sensory Engineering Workshop in Erlangen and the Winter Workshop in Philosophy of Perception in San Diego.

- Dilworth, John (2010), "Realistic Virtual Reality and Perception," *Philosophical Psychology*, 23, 1, pp. 23-42, DOI: [10.1080/09515080903533942](https://doi.org/10.1080/09515080903533942).
- Fulkerson, Matthew (2014), "Explaining Multisensory Experience," in *Consciousness Inside and Out: Phenomenology, Neuroscience, and the Nature of Experience*, ed. by Richard Brown, Springer, Dordrecht, pp. 365-373.
- (2021), "Peripersonal space, bodily self-awareness, and the integrated self," ed. by De Vignemont, Serino, Wong, and Farne, pp. 251-265.
- Hill, Christopher S. (2023), "Perceiving Particulars," in *Contemporary Debates in Philosophy of Mind*, ed. by Brian P. McLaughlin and Jonathan Cohen, Second Edition, Wiley-Blackwell, Hoboken, chap. 24, pp. 433-448.
- (n.d.), "How Should We Understand VR, and What Can It Teach Us?"
- Juul, Jesper (2005), *Half-Real: Video Games between Real Rules and Fictional Worlds*, The MIT Press, Cambridge, MA.
- Kennedy, Chris (2007), "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives," *Linguistics and Philosophy*, 30, 1, pp. 1-45.
- Klein, Ewan (1980), "A Semantics for Positive and Comparative Adjectives," *Linguistics and Philosophy*, 4, 1, pp. 1-45, DOI: [10.1007/bf00351812](https://doi.org/10.1007/bf00351812).
- Kroon, Fred and Alberto Voltolini (2023), "Fictional Entities," in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta and Uri Nodelman, Fall 2023, Metaphysics Research Lab, Stanford University.
- Kvasova, Daria, Laia Garcia-Vernet, and Salvador Soto-Faraco (2019), "Characteristic sounds facilitate object search in real-life scenes," *Frontiers in psychology*, 10, p. 2511, ISSN: 1664-1078.
- Lanier, Jaron (2017), *Encounters with Reality and Virtual Reality*, Henry Holt and Company.
- Lewis, David K. (1978), "Truth in Fiction," *American Philosophical Quarterly*, 15, 1, pp. 37-46.
- Logue, Heather (2023), "Abstract and Particular Perceptual Content: The Best of Both Theories," in *Contemporary Debates in Philosophy of Mind*, ed. by Brian P. McLaughlin and Jonathan Cohen, Second Edition, Wiley-Blackwell, Hoboken, chap. 25, pp. 449-465.
- Loiperdinger, Martin and Bernd Elzer (2004), "Lumiere's Arrival of the Train: Cinema's Founding Myth," *The Moving Image*, 4, 1, pp. 89-118.
- Macpherson, Fiona (2014), "The Structure of Experience, the Nature of the Visual, and Type 2 Blindsight," *Consciousness and Cognition*, 32, pp. 104-128, DOI: [10.1016/j.concog.2014.10.011](https://doi.org/10.1016/j.concog.2014.10.011).
- (2020), "Is Virtual Reality Experience Veridical, Illusory or Hallucinatory? A Complex Answer Based on a New Theory of Illusion and Hallucination and the Nature of the Technology Used to Create Virtual Reality," ms., University of Glasgow, <https://eprints.gla.ac.uk/226457/>.
- Macpherson, Fiona and Clare Batty (2016), "Redefining Illusion and Hallucination in Light of New Cases," *Philosophical Issues*, 26, 1, pp. 263-296, DOI: [10.1111/phis.12086](https://doi.org/10.1111/phis.12086).
- Matthen, Mohan (2005), *Seeing, Doing, and Knowing: A Philosophical Theory of Sense Perception*, Oxford University Press, Oxford.

- McDonnell, Neil and Nathan Wildman (2019), "Virtual Reality: Digital or Fictional?" *Disputatio*, 11, 55, pp. 371-397, DOI: [10.2478/disp-2019-0004](https://doi.org/10.2478/disp-2019-0004).
- McGurk, Harry and John MacDonald (1976), "Hearing lips and seeing voices," *Nature*, 264, 5588, pp. 746-748.
- Meskin, Aaron and Jonathan Cohen (2008), "Photographs as Evidence," in *Photography and Philosophy: Essays on the Pencil of Nature*, ed. by Scott Walden, Wiley-Blackwell, Malden, pp. 70-90.
- Meskin, Aaron and Jon Robson (2012), "Fiction and Fictional Worlds in Videogames," in *The Philosophy of Computer Games*, ed. by J. R. Sageng, T. M. Larsen, and H. Fossheim, Springer, pp. 201-18.
- Morales, Jorge, Axel Bax, and Chaz Firestone (2020), "Sustained Representation of Perspectival Shape," *Proceedings of the National Academy of Sciences of the United States of America*, 117, 26, pp. 14873-14882.
- Morales, Jorge and Chaz Firestone (2023), "Empirical Evidence for Perspectival Similarity," *Psychological Review*, 1, 1, pp. 311-320, DOI: [10.1037/rev0000403](https://doi.org/10.1037/rev0000403).
- Noë, Alva (2004), *Action in Perception*, MIT Press, Cambridge, Massachusetts.
- O'Callaghan, Casey (2019), *A Multisensory Philosophy of Perception*, Oxford University Press, Oxford, United Kingdom.
- Simeone, Adalberto L. (2018), "Substitutional Reality," in *Encyclopedia of Computer Graphics and Games*, ed. by Newton Lee, Springer International Publishing, Cham, pp. 1-4, ISBN: 978-3-319-08234-9, DOI: [10.1007/978-3-319-08234-9\\_254-1](https://doi.org/10.1007/978-3-319-08234-9_254-1), [https://doi.org/10.1007/978-3-319-08234-9\\_254-1](https://doi.org/10.1007/978-3-319-08234-9_254-1).
- Skipper, J I, V van Wassenhove, H C Nusbaum, and S L Small (2007), "Hearing Lips and Seeing Voices: How Cortical Areas Supporting Speech Production Mediate Audiovisual Speech Perception," *Cerebral Cortex*, 17, pp. 2387-2399.
- Sullivan, Jeremiah (2018), "The Untold Story of the World's First Photograph," *AirBNB Magazine* (July 2018), <https://medium.com/airbnbmag/the-untold-story-of-the-worlds-first-photograph-277285f2d78a>.
- Tavinor, Grant (2009), *The Art of Videogames*, Wiley-Blackwell.
- (2019), "On Virtual Transparency," *Journal of Aesthetics and Art Criticism*, 77, 2, pp. 145-156, DOI: [10.1111/jaac.12626](https://doi.org/10.1111/jaac.12626).
- Thomasson, Amie L. (1998), *Fiction and Metaphysics*, Cambridge University Press, New York.
- Velleman, J. David (2013), "Virtual selves," in *Foundations for Moral Relativism*, OpenBook Publishers.
- Walton, Kendall (1984), "Transparent Pictures: On the Nature of Photographic Realism," *Critical Inquiry*, 11, pp. 246-276.