

## A Non-Representational Understanding of Visual Experience

Kaplan Hasanoglu

*Emmanuel College*

This paper argues that various phenomenological considerations support a non-representational causal account of visual experience. This position claims that visual experiences serve as a non-representational causally efficacious medium for the production of beliefs concerning the external world. The arguments are centered on defending a non-representational causal account's understanding of the cognitive significance of visual experience. Among other things, such an account can easily explain the inextricable role that background beliefs and conceptual capacities play in perceptually-based external world belief-formation processes, the fact that visual mental states constrain beliefs because of their presentational phenomenology, and the phenomenon known as the transparency of visual experience.

Keywords: perceptual experience, hallucination, phenomenal properties

In this paper, I will show that certain phenomenological considerations support a *non-representational causal account* of visual experience.<sup>1</sup> This is a position in the philosophy of mind that claims that visual experience serves as a non-representational causally efficacious medium for the production of beliefs concerning the external world. So, it challenges both what Brewer (2006) calls *the content view*, which is the popular philosophical position that claims that visual experience is constituted by representational contents, as well as less fashionable alternatives such

---

Correspondence concerning this article should be addressed to Kaplan Hasanoglu, Ph.D., Philosophy Department, Emmanuel College, 400 The Fenway, ADM 357, Boston, Massachusetts 02115. Email: hasanoglu@emmanuel.edu

<sup>1</sup>Here and throughout I use the term “phenomenological” in the familiar sense that pertains to both *what it is like* to experience the world, as well as any related conclusions. Introspective reports are notoriously unreliable, given the way that they can be easily influenced by, say, folk psychology and other dubious theories. Therefore, in what follows I will methodically and self-consciously attempt to avoid — or, in Husserlian terminology, “bracket” — as many questionable theoretical commitments as possible. I am not alone in thinking that this kind of investigation is both still possible and potentially fruitful — see, for example, Gallagher (2005) and Thompson (2007).

as naïve realist (disjunctivist), sense-data, and adverbialist accounts. My aim is not to argue directly against any of these positions, but rather to present a non-representational causal account as another position worth taking seriously.

In what follows, when I talk of a perceptual experience's "cognitive significance," I will be referring solely to the role it plays in external world belief formation and sustenance, rather than in helping to render some external world beliefs justified and/or reasonable. That perceptual experience has the former role is fairly uncontroversial. For example, at the moment you have various (dispositional) beliefs about your surroundings, and it seems plausible to say that your current perceptual experiences play a crucial role in causing and sustaining those beliefs. A non-representational causal account understands this (roughly) as follows: in veridical cases, a non-representationally constituted internal mental state helps to produce a (reliably) true belief; whereas in non-veridical cases it obviously does not.

As a way of framing my task, let me mention and then somewhat abruptly set aside two possible objections. First, throughout this paper I will be unapologetically assuming the existence of reflectively accessible mental states as constituents of both veridical and hallucinatory experiences. That such states exist has famously come under repeated attack (making my assumption here the source of one objection). However, rightly or wrongly, the philosophical debate between content theorists, naïve realists (disjunctivists), sense-data theorists, adverbialists, etc., presupposes the existence of such states: the disagreement between them surrounds their status in hallucinatory cases, and/or the extent to which a proposed account of them in hallucinatory cases can be extended to the mental states involved in veridical cases. The purpose of this paper is to present a non-representational causal account as another horse in that particular race, so I will not here question the existence of such states. This may turn out to be a problematic theoretical commitment (see footnote 1 above); but given my aim, it is a necessary one.

Next, even if they do exist, it might appear *prima facie* problematic to posit such mental states as a non-representational causal intermediary involved in belief-formation, at least if we are also to grant that perceptual experiences play some role in *justifying* external world beliefs. Natural causes aren't reasons, after all, and in what follows I will only speak of experiences as causes. However, for one thing, though I can't defend the claim here, I agree with Shaun Gallagher that: "To have a belief is not to have an all-or-nothing mental representation, but to have some more-or-less-complete set of dispositions to act and to experience in certain ways" (2005, p. 214). So, in my view, what it means to say that a belief is an intentional entity is, at best, far from clear. If nothing else, this muddles the issue: What exactly is the problem here for a non-representational account supposed to be? And even if we ignore issues surrounding how to understand the intentionality of beliefs, it is also crucial to note that there is no current consensus

in the philosophical literature regarding how to understand *justification*, generally speaking. This matters here because the stand one takes on, say, the internalist vs. externalist and/or the foundationalist vs. coherentist debate will have obvious implications for one's views about the justificatory role of perceptual experience. So-called *dogmatists* like Pryor (2000) and Heumer (2001) are (what I would consider to be) internalists who maintain that perceptual experiences provide prima facie justification for certain external world claims in virtue of their constitutive representational content. By contrast, Davidson (1983) famously defends a coherentist position that maintains that perceptual experiences cause but do not justify external world beliefs. Given such complexities, I have chosen to try to divide and conquer in this paper, and hence will leave off discussing perceptual experience's justificatory role for another occasion. Of course, punting on this issue might *still* seem unfair, since one reason for adopting a content theorist's position is its well-recognized theoretical elegance on this front — constitutive representational content makes the justificatory role of perceptual experience regarding beliefs with the same content a straightforward matter. But such support for a content theorist's view is obviously defeasible, and a complicated meta-methodological question surrounds how to weigh this kind of theoretical support against what I will argue below is recalcitrant phenomenological data. So, in sum, although it is fair to point out that a non-representational causal account owes us an explanation of the relationship between perceptual experience and justified beliefs about the external world, I would argue that it is unfair to rule it out initially based upon the presence of this explanatory gap. Indeed, if nothing else, the phenomenological investigation involved in what follows may turn out to be preparatory for filling that gap.

### *Phenomenology and the Cognitive Significance of Perceptual Experience*

So, let me now clarify my aim. The phenomenological case I make for a non-representational causal account shall be centered, in particular, on defending its understanding of the role of reflectively accessible internal mental states in external world belief-formation, with a focus on visual experience. (I do think the view generalizes to other sense-modalities, but I will not defend that more general claim here.)

To motivate my position, it will help to begin by considering recognized visual perceptual illusions. In such cases, we are easily able to distinguish (i) what it means to take the experience at face value; (ii) what, all things considered, we take to be the case on the basis of the experience; and (iii) the inextricable role background beliefs play in having it be that (i) and (ii) come apart for such cases. Suppose that one is looking at a stick half-submerged in a glass of water. One natural (though not undisputed) thought is that if one were to take that experience at face value, one would take the stick in question to be broken in two. However, because one

has background beliefs about (for example) the physics of light, one does not, all things considered, take the stick to be that way.

For cases like this, the distinction and interplay between (i)-(iii) seems relatively clear. What about more ordinary experiences? To help address this matter, I will now introduce some terminology. Take *the Good Case* to be a visual veridical case that is not an example of a veridical hallucination.<sup>2</sup> By contrast, *the Bad Case* should be understood to be a perfect (non-veridical) hallucination that is indistinguishable through reflection from some related Good Case. Finally, I will reserve the term *ordinary experience* for an experience of a typical human subject with relatively good vision that is either some Good Case involving unobstructed middle-sized objects in the viewed vicinity, or else the related Bad Case.

Consider, now, an ordinary experience of a nearby red sphere. I admit that by endorsing the claim that there is a red sphere, I *seem* to be doing nothing other than endorsing the experience itself, understood as an internal mental state, and in a way that intimately depends on how the mental state phenomenally presents things from my perspective. At the very least, this might seem to support the content view. However, against this, one should first acknowledge that if my background beliefs were to change significantly enough, a contrary claim would be endorsed upon enjoying a phenomenally identical experience. Examples that illustrate this are familiar: depending on my background beliefs, on the basis of the same internal visual state (here and henceforth understood as a structured complex of phenomenal properties), I might believe that I'm looking at a red wall, or else one that is white but bathed in red light. According to a non-representational causal account, the best explanation for this is that there is no one set of accuracy conditions concerning the external world that an experience can intrinsically call its own (see Travis, 2004, for an argument against the content view that seems to rest on a similar point).

What might be a typical example of such a background belief for ordinary experiences? Well, it is obviously part of my conception of spatiotemporal objects that they remain relatively stable through time: that is, we don't think of a mind-independent object as the kind of thing that will wink out of existence once we stop looking at and/or interacting with it. In other words, the size, shape, etc. of an object is what I will henceforth call an *experientially transcendent property*. Generally speaking, a property of an object  $x$  is an experientially transcendent property at time  $t$  — where  $t$  represents a time period during which we experience  $x$  — when, *ceteris paribus*,  $x$  also has that property at times other than  $t$ . Mind-independent objects are relatively stable, then, because their size, shape, etc. are experientially

---

<sup>2</sup>As the label suggests, veridical hallucinations are hallucinations that happen to be accurate, even though the cause of the perceptual experience is not the relevant object that is in fact in the perceiver's surroundings. For (a well-worn) example, I may hallucinate a dagger before me while there just so happens to be a dagger of that very same sort before me. For useful discussion, see Johnston (2004). For discussion of the related notion of veridical illusions, see Johnston (2006).

transcendent properties. According to a non-representational causal account, for cognitively sophisticated adults, beliefs about experientially transcendent properties are thus the kind of background beliefs that help explain why we believe what we do when we experience things like red spheres. On this view, the visual state does not itself present a red sphere, and it is, rather, background beliefs such as those concerning experientially transcendent properties that pick up the slack in the relevant belief-formation processes. Perceptual experience, properly construed, ought to be relegated to a mere non-representational causal role in that process.<sup>3</sup>

However, rival positions such as the content view can certainly accommodate the role that background beliefs play in such belief-formation processes.<sup>4</sup> Moreover, one might think that a non-representational causal account should be rejected on straightforward phenomenological grounds, since it seems unable to explain the way that perceptual mental states constrain belief. So, for example, it is fairly obvious that the mental state I'm enjoying at the moment constrains my beliefs in many ways — for one thing, it prevents me from believing that I'm walking down a busy street. Crucially, it evidently does this because of its *presentational phenomenology* — the mind-independent world is presented by the mental state I'm enjoying right now as being a way that my current (dispositional) beliefs reflect, and this feature seems intrinsic to it. (For discussion of this, see Siegel, 2010, chapter 2). Obviously, then, in order to defend it I need to provide an explanation of how internal mental states constrain beliefs on behalf of a non-representational causal account.

To do so, I will now argue that for a typical human subject who regularly enjoys Good Cases, the mental state involved in the Good Cases — to introduce one last piece of terminology — is *specific-object-involving*. A visual state is specific-object-involving when and only when one enjoys it while causally interacting in a non-deviant way with the mind-independent object(s) that play(s) an essential role in helping to create it. The idea here is that, since it is a Good Case, the mind-independent objects that are actually before you play an essential and non-deviant role in helping to create your current token visual state, similar to how the actions of a painter

---

<sup>3</sup>By contrast, in an influential paper, Horgan and Tienson (2002) defend a position that evidently would involve denying that background beliefs are relevant as to why we take certain experiences to be of the external world. They maintain, or so it appears, that intentional content involving the external world is intrinsic to the phenomenal properties of the experience of, say, a red sphere in one's immediate vicinity. But, as far as I can tell, the force of their arguments rests on introspection alone, and hence can be straightforwardly undermined by the considerations offered throughout this paper. In short, I would argue that Horgan and Tienson simply gloss over the implications of the obvious role that background beliefs, capacities, etc. play in mediating the relevant belief-formation processes, and hence in helping to explain the cognitive significance of the relevant perceptual states.

<sup>4</sup>To defend their view in light of the above kind of holism, the proponents of the content view might draw an analogy with belief. After all, even if one belief by itself doesn't fix what will be justified and/or endorsed upon maintaining it, this alone need not mean that the belief doesn't have determinate content.

are essential to helping to create a masterpiece. So, the visual state you are enjoying at this moment is specific-object-involving, since it exists now only because the objects right in front of you that have helped to create it also exist.<sup>5</sup> This is of course not to deny that such a mental state might have been created some other way. Instead, it is a non-counterfactual claim about what is in fact essential for creating your current token mental state.

I will now spend a fair bit of time offering various qualifications and developments of this idea, including laying out some of its more important implications. Once all of that is in place, I will be in position to provide a non-representational causal account's understanding of the way in which visual experiences constrain beliefs. First, the claim here should not be confused with a claim about what *constitutes* the mental state. To give a helpful analogy, consider that your parents played an essential role in your creation. But although (more specifically) their past actions were essential to creating you, neither your parents nor their actions are literally a part of you. In like fashion, on this view, although they are essential in helping to create them, mind-independent objects should not be thought of as constitutive of the internal visual states enjoyed in Good Cases.

Next, maintaining this is quite compatible with the so-called *abstractness* of perception noted, for example, in Tye (1995) — the fact that in other Good Cases numerically distinct objects can help to produce numerically distinct mental states of ostensibly the same phenomenal type. Since each *token* mental state's phenomenal properties are created, in part, by the mind-independent objects present, all that follows is that any identity of phenomenal *type* will itself be determined, albeit only in part, by the similarities of the distinct, specifically perceived, mind-independent objects themselves. So, right now the experience I'm having is specific-object-involving simply because this particular computer that I am now interacting with is playing an essential role in creating its token phenomenal properties. In other words, my current visual state is *this-computer-involving*. All that these latest considerations require me to go on to admit, then, is that the computer in front of me thereby creates a token mental state whose intrinsic phenomenal type is identical with the type of any numerically distinct mental state produced, in part, by the perception of any other sufficiently similar computer.

If possible, as a phenomenological exercise, it would be helpful to now examine one's own visual experience and try this last sort of assertion on for size ("My current visual experience is *this-X-involving*"; where "X" refers to some unobstructed middle-sized object in the viewed vicinity). My prediction is that it will strike one as a very natural assertion to make. On the other hand — to now lay my cards on the table — if the reader finds that there is absolutely nothing to such a characterization, then the arguments that follow will have little force.

---

<sup>5</sup>Here and throughout I have assumed that there is no causal over-determination involved in what in fact creates the mental state. There would be if, for example, it was a veridical hallucination.

Next, and quite crucially, matters are not all that different when it comes to the mental state involved in the Bad Case. You are probably not vividly hallucinating anything right now, but on some other occasion you might. How does a non-representational causal account treat such cases? In a word, that mental state is *derivatively* specific-object-involving. To get at what this means, notice first that in order for the mental state in question to be a perfect hallucination it has to *seem real*; that is, it has to be an experience that is taken to be veridical by the subject suffering from it at the time in question. But that means that *for those who regularly enjoy Good Cases*, the hallucination has to have the same general sort of phenomenal properties as the mental states that were, in fact, created in part by the presence of mind-independent objects the subject has in fact already encountered. It is only a perfect hallucination because it has an intrinsic phenomenal character of a basic sort already possessed by the mental states involved in the already experienced Good Cases. Intuitively, for those who regularly enjoy Good Cases, it is hard to understand how one could have a perfect hallucination — one indiscriminable through reflection from some related Good Case — that instantiates intrinsic phenomenal properties substantially different from the basic type of phenomenal properties created by Good Cases on other, previous occasions. Suppose, then, I am suffering from a perfect hallucination of a computer. In that case, the mental state can't be this-computer-involving in the sense mentioned above, simply because I am not causally interacting with any computer. However, since the hallucination would seem real, whatever caused it to come about would have to generate phenomenology that copied my normal, everyday experiences of the external world, including things like computers. It is in that sense that the mental state in question would be derivatively specific-object-involving.

To clarify, perhaps another analogy can help. For, although it shares observable properties with the original, you can only understand a perfect forgery of some painting in a similar manner. Dali painted *The Persistence of Memory*, but he did not paint a certain perfect forgery of it. And yet, by virtue of its observable properties, a certain painting is in fact a perfect forgery of *The Persistence of Memory* because of its derivative relationship with Dali's original. That is, we have to reference Dali's actions with respect to the original, if we are to understand why it is correct to describe a painting with certain observable properties as a perfect forgery of *The Persistence of Memory*, rather than as (of) some other painting. With important qualifications, hallucinations can be understood similarly. Admittedly, no analogous ontological priority is involved here. That is, we arguably don't need to first experience an object in order to hallucinate something involving it — fanciful dreams of fanciful things are common enough. So, plausibly, we can have perfect hallucinations involving objects we've never actually seen before. And even in cases where the requirement for prior perception would make sense — say, a vivid dream involving one of my brothers — the situation hallucinated may be novel. I may dream of Taner doing something I've never seen



him do before, wearing clothes I've never seen him wear before, etc. Similarly, I may dream of a banana or another commonly perceived object as seen from a novel angle, in weird lighting, etc. Nevertheless, perfect hallucinations of novel situations and/or of never before seen objects are also derivatively specific-object-involving. No mere smattering of paint on a canvas is a forgery; similarly, no mere play of sensations is a perfect hallucination. In order for a painting to be a perfect forgery, it must be a convincing reproduction. Similarly, in order for an experience to be a perfect hallucination, the subject must take it to be veridical at the time she experiences it. *For subjects who regularly enjoy Good Cases*, this means that in order to understand what makes an experience a perfect hallucination, we must first understand the basic phenomenal character of the mental states involved in their prior everyday veridical experiences. So, if I am going to have a perfect hallucination of a banana (or else of some object I've never actually seen before), then even if the situation, perspective, etc. is novel, it still must be a mental state that shares a basic phenomenal character with the mental states involved in the actual situations I *have* already experienced. In particular, the hallucinated object(s) must still seem to be behaving in generally believable ways. For example, you can't have a perfect hallucination of objects that, *ceteris paribus*, unpredictably flick in and out of existence, spontaneously melt into one another, etc. If I was suffering from a hallucination where things were too weird or otherwise off in some fundamental sense, then assuming I was in an otherwise normal state of mind, I wouldn't take the experience to be veridical. Instead I would probably think something like: "This is too weird, I must be hallucinating." Therefore, it would not be a perfect hallucination. In sum, for those who regularly enjoy Good Cases, it is the prior experiences they have had of actual enviroing mind-independent objects that first determine what it means for things to seem real or not. As a result, it will be its fundamental relationship to those same prior experiences which will determine whether or not my novel dream of a banana, my brother, or even some object I've never seen before, etc., counts as a perfect hallucination. That is what makes even perfect hallucinations involving novel situations and/or never before seen objects derivatively specific-object-involving.

Interestingly, this all relates in important ways to the phenomenon that is often called the *transparency* or *diaphanousness* of perceptual experience, which is something usually thought to support a certain version of the content view. (For discussions of transparency, see Dretske, 1995, p. 62; Harman, 1990, p. 39; Moore, 1903, p. 450; and Tye, 2000, pp. 51-52). Some philosophers have argued that when we reflect on the subjective properties of our perceptual experiences, all we seem to find is what our experiences present as being the case in the mind-independent world. To focus on what it is to subjectively experience blue, for example, is just to focus on the blue thing that one's experience is presenting as being in front of one. What this is typically thought to show is the non-existence of non-representational phenomenal properties; that instead our experiences



only have (sense-modality-specific) representational properties (see Harman, 1990; Tye, 1995, 2000). But one might also employ transparency as a means for claiming, on phenomenological grounds, something more general: that the experience itself has representational content that bears upon the external world. The claim would be that were we to try to find non-representational features of our experiences (such as their non-representational causal features), all we would end up noticing is how they present mind-independent reality.

However, even granting that perceptual experience is transparent, an entirely separate question is how we should interpret this phenomenon (see Stoljar, 2004). Indeed, the proponent of a non-representational causal account can explain transparency as follows. In the Good Case or its related Bad Case, when we examine our internal mental states we only find the world because that mental state in question is one that makes things seem real, where *for those who have regularly experienced Good Cases* this means that the mental state in question has the general sort of non-representational phenomenal properties created by the relevant mind-independent objects one has already experienced. On this view, then, transparency (along with the property of seeming real) is not something intrinsic to an experience, but is rather determined (in part) by the prior environmental objects encountered as one has more or less successfully made one's way through the world.

This allows me to respond to another possible objection. I have in mind a worry that derives from the recent work of William Fish (2009). Fish is a naïve realist. He maintains that the phenomenal character of veridical experience is constituted by a subject's acquaintance with the properties of the experienced object(s) [Fish, 2009, p. 14]. So, for example, when I have a veridical experience of a red tomato, Fish would claim that the phenomenal redness that constitutes the mental state is a property of the tomato that is actually before me. Good Cases thus involve an "irreducible" mental relation with certain mind-independent objects (Fish, 2009, p.14, n.19). Of course, perfect hallucinations would involve no such relation. To account for them, Fish thus argues rather strikingly that perfect hallucinations lack phenomenal character altogether (p. 93). The reason Fish makes this bold move, it seems, is because on his view phenomenal character is something constituted by the above-described, irreducible mental relation. And so, the reasoning seems to go, since that mental relation is lacking in perfect hallucinations, phenomenal character must also be lacking (Schellenberg, 2013, p. 50). Fish admits that perfect hallucinations seem real. But, he argues, this is because of their "cognitive effects," rather than their phenomenal character (p. 94). In particular, they seem real only because they produce "the same beliefs or judgments that a veridical perception of that kind would have produced" (p. 94). For this reason, he would obviously deny that perfect hallucinations seem real because of their derivative phenomenal character.

Fish's fully developed views are complex, and his arguments are characteristically sophisticated and subtle. They are also stated within the context of the rather

involved debate over naïve realism. Suffice it to say, then, that I cannot give Fish a fair treatment here. In any case, there is a simple and glaring problem for him. As Susanna Siegel bluntly puts it in one of her responses to Fish: “The idea that hallucinations lack phenomenal character is at odds with the crudest deliverances of introspection” (Siegel, 2010, p. 49, n. 19; for a more detailed critique, see Siegel, 2008. See also Martin, 2013; Pautz, 2013; and Schellenberg, 2013). As I would put it, a very basic phenomenological investigation supports the claim that perfect hallucinations possess phenomenal character. As a result, one would need extremely compelling reasons for denying that hallucinations actually possess such character.

I don’t think that Fish provides such reasons. I have argued that Good Cases are specific-object-involving, but in a sense that amounts to a non-deviant *causal relationship* rather than a constitutive and “irreducible” relationship of acquaintance between the mental state and the object(s) before one. I am no naïve realist. And *if*, in order to maintain naïve realism, one must also deny that perfect hallucinations have phenomenal character, then I would argue that that is too high of a price to pay. (Whether this conditional actually holds is not something I can properly address here.) In sum, arguably, my position is superior to Fish’s in the following way: I can explain a sense in which the mental states in Good Cases are world-involving, while also acknowledging that perfect hallucinations possess phenomenal character.

There is one final important implication of my overall position that is worth discussing at length. To get at it, we should note that one ubiquitous example of an object involved in Good Cases will be one’s physical body — though, of course, one’s body is special in that it is an object that one *lives through* rather than experiences at a distance. I mention this here because, since one’s lived physical body is a unique item situated in the universe, a non-representational causal account is thus nicely poised to provide a way of *individuating* visual mental states. To individuate something is to provide an adequate account of what makes it unique. As I will understand it, such an account must have an actual and a counterfactual component.<sup>6</sup> It must both provide a true description of just the mental state in question *and* also fail to allow for (nearby?) possible worlds where there is more than one mental state that satisfies that same description (though there can be possible worlds where a different mental state satisfies the description). The first requirement assures that the properties mentioned in the individuating account are uniquely true of the mental state in question; the second assures — to a degree arguably suitable for my purposes — that the fact that those properties *uniquely* pick it out in the actual world is not an accident. (There are potentially going to be many different ways to individuate mental states in this way.)

---

<sup>6</sup>I am not sure how to understand individuation, generally speaking. However, here I am only concerned with individuating mental states adequately enough for the purposes of understanding the cognitive significance of perceptual experience. For that project, I would argue that the account offered here suffices.

To consider a different sort of case, on this view one adequate way to individuate Barack Obama is as 44<sup>th</sup> President of the United States. This is because he is in fact the only one that held that office *and* there are no (nearby?) possible worlds where there is more than one 44<sup>th</sup> President of the United States, given how we now understand that institution. (And, as predicted, there are other ways to individuate Barack Obama. For example, one might cite specific enough aspects of his unique personal history.)

So, as far as any visual states enjoyed in Good Cases are concerned, my claim here is that one way to individuate them involves citing the mind-independent objects (including one's body) that have helped to create them. To get at *why*, it will be useful to explore some related yet inadequate attempts to individuate such states. For example, notice that if you were to merely focus on a certain visual state's intrinsic phenomenal properties as a way of individuating it, mentioning nothing about its etiology, the account would be too coarse-grained: you would here run the risk of wrongly identifying it with a numerically distinct but phenomenally type-identical experience that is or might be enjoyed by some subject (including oneself on a different occasion). In other words, there is a possible world (that is perhaps the actual world) where there is more than one visual state with those same phenomenal properties. Merely adding the actual time that you enjoyed the state runs into similar problems: here you would run the risk of misidentifying it with a numerically distinct phenomenally type-identical experience that is or might be enjoyed by someone else at that particular time. That is, there is a possible world (that may be the actual world) where there is more than one visual state with those phenomenal properties enjoyed at that time. This suggests that one way to make the account sufficiently fine-grained is to add an adequately specified *place* where it is enjoyed. So, you might try to home in on the subject's unique place and time where he enjoys the state. My claim, then, is that one way of adequately specifying such a place and time is by citing the specific enviroing objects that have helped to create the mental state, including their precise location in relation to the subject's body. As long as this citation is specific enough in this way, referring to the specific object(s) rather than object type(s), it seems it will allow us to adequately individuate the mental state in question. There is no one other than myself (the body that is currently enjoying this visual state) viewing this particular computer from this particular angle relative to my body. And it certainly stretches the imagination to consider a case where there is more than one such visual state enjoyed at a certain time by an embodied subject like myself. Roughly, then, I would argue that an adequate principle of individuation for visual states enjoyed in Good Cases seems to be this: at any time *t*, visual states *x* and *y* are identical if and only if they have the same intrinsic phenomenal properties and the same mind-independent etiology (where we suitably specify that etiology according to one's unique bodily-oriented perspective). Notice, then, that on this view subjects involved in Good Cases can individuate their own mental states by making use of this principle, through a fairly unsophisticated phenomenological exercise like the one rehearsed a few sentences ago.

Typically, in discussions of perceptual experience, mind-independent objects are thought to be only extrinsically related to internal mental states, as an aspect of their causal ancestry. The object is thus not thought of as something given in the experience itself, since the latter is understood as a relatively independent downstream event in the overall process. This common way of looking at visual experience thus treats a mental state enjoyed in a Good Case, considered as such, as an ontological (and perhaps also epistemological) par with a phenomenally indistinguishable hallucination. It is a way of looking at perceptual experience that is also friendly to the methodology enshrined by the Argument from Hallucination, where one first examines perfect hallucinatory cases and then attempts to spread that understanding to veridical cases, as well. However, according to the account defended here, we cannot rely on this outlook or this methodology if we wish to understand the cognitive significance of perceptual experience. If that is our aim, then we must acknowledge that in Good Cases visual mental states are merely one component of a larger process; a process that also involves background beliefs *and* the mind-independent objects (including one's physical body) partly responsible for creating the state. To say that mind-independent objects are only extrinsically related to these internal mental states is thus misleading, since it ignores the absolutely essential role that such objects play in helping to create the token mental state involved in a Good Case, and also the role they can play in such a case in an attempt to individuate one's own visual states through a phenomenological exercise. It also, for that matter, fails to take into account the derivative status of the mental states involved in Bad Cases. On the contrary, then, it is not despite but *because* a mind-independent object (including one's body) plays the role that it does in an internal mental state's causal ancestry in Good Cases, that we must include the former as essential to understanding the latter. So, since my current mental state is this-computer-involving, it follows that it is the position of the computer that is actually in front of my physical body at the moment that helps to determine the phenomenal properties of my current internal mental state; moreover, to the extent that this computer's position is something I am able to control through bodily action, that mental state likewise determines the position of the computer. To deny the computer's essential role in making my current mental state what it is, therefore, would be just as absurd as denying that my parents played an essential role in creating me, or that a painter played an essential role in creating a masterpiece, etc. Even though mind-independent objects are not, on this view, literally parts of the visual states that they help to create, since they are essential to their creation and also something we can cite to individuate such states via a phenomenological exercise, we cannot and should not ignore them when attempting to understand the latter; or, at least, we cannot and should not if our goal is to understand the cognitive significance of visual experience.

The phenomenological linchpin that brings these ideas home is this: when you are involved in a Good Case you simply cannot refer to an experience *as some isolated internal mental event or state* except by way of a mental act which

*abstracts from* the lived, contextualized, body-and-world-involving process that is the Good Case as a whole, understood as a kind of successful dance with various immediately enviring objects. According to a non-representational causal account, an act of abstraction like this is precisely what I engaged in above, when I recognized that my current mental state is this-computer-involving.

Now for the punchline: according to a non-representational causal account, the mental states involved in veridical experiences and perfect hallucinations constrain beliefs in the way they do simply because those mental states are specific-object-involving, and derivatively specific-object-involving, respectively. As a result, whatever specific mind-independent object is involved in helping to create a token mental state of a certain type (understood as a structured complex of phenomenal properties) will obviously thereby determine what beliefs result, albeit relative to a fixed set of background beliefs, capacities, etc. The mental state that is (derivatively) this-computer-involving will thereby not produce a belief that I am looking at a lion, relative to a fixed set of background beliefs, capacities, etc., simply because computers and lions are very different kinds of objects that consequently thereby help to produce very different structured complexes of phenomenal properties. In the Good Cases involving computers or lions, the different way in which such mental states constrain my beliefs, therefore, rests on nothing other than the mind-independent differences between computers and lions. The same is true in the Bad Cases derivatively involving computers or lions. In other words, a hallucination that is derivatively lion-involving will not make me believe that a computer is present relative to my current fixed set of background beliefs, capacities, etc. simply because the lions that I have already experienced have been different enough in a mind-independent sense from the computers I have already experienced, and hence have helped to produce the appropriately different structured complexes of phenomenal properties that thereby help to produce the appropriately different beliefs.

However, it must be acknowledged that being an object such as a computer is not simply a matter of having, say, a certain mind-independent molecular structure, but is also a function of the object in question's (socially reinforced) role in our practices. So, if we just so happened to give a certain mind-independent swarm of atoms a different role in our practices (say, as an object of worship rather than as a computing device), it would of course no longer be a computer. One might say, then, that the computer before me is only a computer because we represent it as a computer.

A similar point has been forcefully made by McDowell, who famously uses it to defend a certain version of the content view. In *Mind and World* he argues that we only experience the world as we do because we already possess and bring to bear various conceptual capacities. So, on his view, we can experience something as a computer only because we have such capacities already in place; capacities which (as I understand McDowell) have already holistically situated what it means to be a computer in a way that allows for the experience to rationally link up with other beliefs, experiences, actions, etc. As McDowell writes:

By virtue of the way in which the conceptual capacities that are drawn into operation in an experience are rationally linked into the whole [conceptual] network... the subject of the experience understands what the experience takes in (or at least seems to take in) as part of a wider reality, a reality that is all embraceable in thought but not all available to this experience. (1996, pp. 31–32)

However, although it is quite correct that being a computer is in part a function of this kind of representation (and hence of what conceptual capacities we bring to bear when we visually perceive the world), this can be straightforwardly handled by a non-representational account in the manner already discussed. Namely, one can claim that what is already in place and hence thereby contributes to that representing is itself a function of our background beliefs, conceptual capacities, etc. rather than something intrinsic to the mental state.

To reinforce this last point and connect it up with the latest worry, suppose that I gain the concept of a Douglas fir.<sup>7</sup> Then suppose I subsequently see one in a set of circumstances (lighting, level of sobriety, etc.) that are otherwise identical with some set of prior circumstances in which I was in the presence of the same tree, but where I didn't possess the concept. Suppose, next, that (partly) on the basis of the latest experience I believe that a Douglas fir is present, but that I didn't do so on the basis of the previous experience. There are two possibilities here: either the later mental state's intrinsic phenomenal properties are type-identical with those that constituted the earlier mental state, or they are not. Consider the first option. In that case, one seems free to maintain that it is the very same type of perceptual mental state involved in both cases. (Why can't we slice things up in this way?) The proponents of a non-representational causal account seem free, then, to give that common type of mental state a non-representational gloss. They could maintain that since the intrinsic phenomenal properties did not change, the type of perceptual mental state involved in the relevant belief-formation process also did not change, even after I gained the concept of a Douglas fir. Instead, it is only the relevant background beliefs, capacities, etc. that are now different. According to a non-representational causal account, then, it would be the latter difference alone which would explain why I now believe that a Douglas fir is present, whereas before I didn't.

Similarly, even if for the sake of argument we granted that the later experience of the Douglas fir was constituted by different phenomenal properties, it still would be far too hasty to maintain that that change in particular is what fundamentally explains the difference in cognitive significance. Given that background

---

<sup>7</sup>This example could obviously be adapted to apply to the computer vs. object of worship case, *mutatis mutandis*. In particular, rather than talking about gaining a concept, we could talk about an alteration of the relevant conceptual capacity that I bring to bear in perceiving the swarm of atoms now identified as a computer.

beliefs, capacities, etc. inextricably mediate the belief-formation process, another possibility is that as before it is only a change in such beliefs, capacities, etc. that explains why I now believe that a Douglas fir is present, whereas before I didn't. In other words, it may well be that the change in phenomenal properties is simply explanatorily irrelevant.

More generally speaking, there is little doubt that changing what we believe and how we think (including the gaining/losing of conceptual capacities and related dispositions) can change how things perceptually seem to us — that there is so-called *cognitive penetration* involved here. But, if nothing else, there is still a debate to be had regarding what precise role such cognitive penetration should play in any attempt to explain the cognitive significance of perceptual experience.<sup>8</sup>

Finally, a related point that should be acknowledged is that when we see a computer before us, we do so via a sense modality that makes its own causal contribution to the process. As a result, in Good Cases, *how* a computer appears to us will, in part, be a function of peculiarities related to how our brain and eyes work. But it still can be true in Good Cases that *what* we see are mind-independent items before us that themselves have relational properties like *looking a certain way to a certain subject in a certain circumstance*, and that these relational properties of that mind-independent object *also* help to determine how we causally interact with it in a non-deviant fashion via that particular sense modality. So, when in the Good Case I see and point at a computer that is actually in front of me and say something like “This computer looks (is) black,” I am pointing at a mind-independent object that looks the way it does in part because of how my brain and eyes work. And I call it a *computer* because of the particular role that that swarm of atoms plays in our practices. But because it is a mind-independent swarm of atoms that looks that way and plays that role (rather than something mind-dependent, like a dream or an after-image), it is nothing other than a mind-independent object which also helps to produce my current internal mental state.

---

<sup>8</sup>Recent work by Fish and Johnston also plausibly accommodates the role that conceptual capacities play in shaping the presentational phenomenology of perceptual experience, without espousing the existence of conceptual content. Reviewing the details here would be rather involved, and thus would perhaps take us too far afield. In any case, if nothing else the prominence of Fish's and Johnston's work alone shows that there is at least a debate to be had here, and hence that the above, McDowell-inspired objection is not decisive. See Fish (2009, pp. 67–74) and Johnston (2006, pp. 282–285). The point I am making above also helps to illustrate, I think, why the Burgean externalist views on mental content cannot be of assistance to the proponent of the content view. After all, Burge's thought experiments can be construed as supporting the idea that one's conceptual capacities are partly constituted by one's environment. But, one could of course admit that and just say that what the external environment thereby helps to constitute is the conceptual capacities that enter into the relevant belief-formation processes, rather than the content of the experience as such. Unfortunately, I lack the space to address the argument for perceptual content found in Burge (2010).



## Conclusion

My hope is that the above discussion has shown how a non-representational causal account can readily explain various aspects of our visual experiences. However, this does not by itself show a non-representational causal account to be correct, or even the most plausible view to maintain. Making that case would require showing it to be the *best* explanation of such phenomena, preferably along with other basic features of perceptual experience. Although I have intimated at times how it might proceed, that project must be reserved for another occasion.

## References

- Brewer, B. (2006). Perception and content. *The European Journal of Philosophy*, 14, 165–181.
- Burge, T. (2010). *Origins of objectivity*. New York: Oxford University Press.
- Davidson, D. (1983). A coherence theory of truth and knowledge. In D. Henrich (Ed.), *Kant oder Hegel?* (pp. 423–438). Stuttgart: Klett–Cotta.
- Dretske, F. (1995). *Naturalizing the mind*. Cambridge, Massachusetts: MIT Press.
- Fish, W. (2009). *Perception, hallucination, and illusion*. New York: Oxford University Press.
- Gallagher, S. (2005). *How the body shapes the mind*. Oxford: Oxford University Press.
- Harman, G. (1990). The intrinsic quality of experience. *Philosophical Perspectives*, 4, 31–52.
- Huemer, M. (2001). *Skepticism and the veil of perception*. Lanham, Maryland: Rowman and Littlefield.
- Horgan, T., and Tienson, J. (2002). The intentionality of phenomenology and the phenomenology of intentionality. In D. Chalmers (Ed.), *Philosophy of mind: Classical and contemporary readings* (pp. 520–533). New York: Oxford University Press.
- Johnston, M. (2004). The obscure object of hallucination. *Philosophical Studies*, 120, 113–183.
- Johnston, M. (2006). Better than mere knowledge? The function of sensory awareness. In T. S. Gendler and J. Hawthorne (Eds.), *Perceptual experience* (pp. 260–290). Oxford: Oxford University Press.
- Martin, M. G. F. (2013). Shibboleth: Some comments on William Fish's *Perception, Hallucination & Illusion*. *Philosophical Studies*, 163, 37–48.
- McDowell, J. (1996). *Mind and world*. Cambridge, Massachusetts: Harvard University Press.
- Moore, G. E. (1903). The refutation of idealism. *Mind*, 12, 433–453.
- Pautz, A. (2013). Do the benefits of naïve realism outweigh the costs? Comments on Fish, *Perception, Hallucination and Illusion*. *Philosophical Studies*, 163, 25–36.
- Pryor, J. (2000). The skeptic and the dogmatist. *Nous*, 34, 517–549.
- Schellenberg, S. (2013). Perceptual content and relations. *Philosophical Studies*, 163, 49–55.
- Siegel, S. (2008). The epistemic conception of hallucination. In A. Haddock and F. Macpherson (Eds.), *Disjunctivism: Perception, action and knowledge* (pp. 205–224). Oxford: Oxford University Press.
- Siegel, S. (2010). *The contents of visual experience*. Oxford: Oxford University Press.
- Stoljar, D. (2004). The argument from diaphanousness. In M. Escurria, R. Stainton, and C. Viger (Eds.), *Language, mind, and world: Special issue of the Canadian Journal of Philosophy* (pp. 341–390). Edmonton, Canada: University of Alberta Press.
- Thompson, E. (2007). *Mind in life: Biology, phenomenology, and the sciences of the mind*. Cambridge, Massachusetts: Belknap Press.
- Travis, C. (2004). The silence of the senses. *Mind*, 113, 57–94.
- Tye, M. (1995). *Ten problems of consciousness*. Cambridge, Massachusetts: MIT Press.
- Tye, M. (2000). *Consciousness, color, and content*. Cambridge, Massachusetts: MIT Press.