

Nature's Providence: The Representational Role of Vision

Tim Klaassen

Abstract

In this paper, I propose that visual perception is a straightforward case of (mental) representation. The phenomenology of vision is key here: as we see, we are directly presented with aspects of the environment that are at various distances away from us. Through the process of vision, aspects of the environment that would otherwise still be unavailable or “absent”, are made (quasi-)available, or (quasi-)present. This already by itself makes vision deserving of the name ‘representational’. Moreover, all of this holds true, even if we suppose that vision is “direct” and intrinsically “contentless”. In part I, I provide a bit of an overview of the relevant theoretical landscape, and I give my own take on how we may frame and approach some of the issues with respect to mental representation. Then, in part II, I argue that the representational features of vision must primarily be sought at the level of conscious experience. In part III, I elaborate on the two most common theories that are associated with the view that vision is ‘representational’, and proceed to show that there is in fact a third, more direct, way in which vision can be considered as essentially a process of representation. Finally, in part IV, I show that the kind of interaction problems that arise with the more traditional notions of representation also arise with this newer, more direct, notion.

I

Why do human beings behave the way they do? According to common sense, people do the things they do because they have certain *beliefs* and *desires*. For example, we might catch a person walking a very intricate route through a city. Upon asking why (s)he is walking this intricate path, (s)he answers that (s)he does so, because (s)he desires to buy some bread, and because (s)he believes that there is a store located at a particular place, and that so and so is the way to get there.

In most of the literature on “folk-psychology”, beliefs and desires are thought of as instances of what are called ‘propositional attitudes’. They are instances in which an agent bears a certain relationship, a certain psychological attitude, towards a proposition. Propositions, in turn, may, for now, be understood as items that specify ‘conditions of satisfaction’.¹ This property is also inherited by the propositional attitudes themselves. To illustrate, the belief that Paris is located about 750 kilometers south of Amsterdam can be *true* or *false*. And the desire for more bicycle sheds in the center of Amsterdam can be *satisfied* or *frustrated*. Propositional attitudes explain human behavior because people are taken generally to behave in accordance with how they take the world to be, and in accordance with how they would *like* the world to be.

We might say that according to “folk psychology”, propositional attitudes are causally responsible for behavior. Whether this is also *true* is, however, a point of contention, especially among philosophers and other theorists of the mind. Indeed, the behaviorists of the beginning of the twentieth century are notorious for having denied the idea that beliefs and desires are among the actual causes of human behavior. One of the more obvious difficulties with the view that human behavior is influenced by beliefs and desires, is that it is committed to the existence of special sorts of *mental* items that, moreover, *causally interact* with the physical world. From the scientific-materialist

¹The exact nature of propositions is a controversial issue among philosophers. But the characterization that I have just given will do for our present purposes.

point of view, such commitments are, on the face of it, problematic.

First of all, mental items, such as beliefs, desires, and other propositional attitudes, don't seem to admit of being scientifically measured. Peer inside the brain for as long as you like; it seems that nothing like a belief or a desire shows up.² Second, many theorists think that the supposition that some physical events have mental causes, is difficult to square with the so-called "causal closure" of the physical world – which is the view that every physical event has a sufficient physical cause.

However, in the fifties and sixties of the twentieth century, a new way of thinking about the mind developed that, some thought, showed that a scientifically and materialistically acceptable account of the folk psychological theory of behavior, might be given after all. This new way of thinking was strongly inspired by the advent of the computer, and has often been referred to as 'the computational theory of mind'. However, because different theorists now seem to have different views about just what computation *is*, this label might be somewhat misleading, because, to put it briefly, it may suggest, to the reader, an inappropriately wide spectrum of theories.³ Therefore, I am going to refer to it, following Andy Clark (Clark; 1989), as '*classical cognitivism*'.

Classical cognitivists maintain that the brain is an organ that is systematically engaged in processes of *formal symbol manipulation*.⁴ What is important about this idea is, among other things, that it supposes that there are items – or patterns of neuronal activation if you will – inside the brain that possess semantic properties. Items that, in other words, carry *meaning*. Moreover, because these items are *physical* in nature, they have *causal powers*. Indeed, according to classical cognitivism, these symbols in fact play a systematic *causal role* in the production of behavior.

²But then again, what do beliefs and desires look like anyway?

³See (Copeland; 1996), (Fresco; 2010), and the various articles written by Gualtiero Piccinini (Piccinini; 2008), (Piccinini; 2009), and (Piccinini; 2012).

⁴For early formulations of this idea, see (McCulloch and Pitts; 1943) and (Newell and Simon; 1976).

The idea that the brain is a formal symbol manipulator, has led some theorist to suspect that there may be a way to show, *how*, and indeed scientifically justify *that*, folk psychological explanations can be true of physical organisms such as human beings.⁵ After all, if the brain is a formal symbol manipulator, then this seems to provide a way of “materializing” the so-called propositional attitudes, and to give them a place and causal role within the physical world. If the propositions that figure in beliefs and desires are expressed symbolically inside the brain, and if the attitudes towards them correspond to specific algorithmic procedures, then, it seems, we are already well on our way of having a scientifically and materialistically acceptable account of folk psychology.

Over the last decades, however, many have, in various ways, expressed their dissent with respect to this promise of a scientific vindication of folk psychology. Some have argued that folk psychology is a false framework to begin with, and that we therefore simply need not bother to try to scientifically vindicate it.⁶ Others argue that, whereas there might be a (greater or lesser) truth in folk psychology, classical cognitivism simply isn’t the right kind of theory to back it up, while the hypothesis that the brain performs, at some level or other, formal symbol manipulations could, despite that, nevertheless be a scientifically fruitful idea.⁷ Yet others, most notably connectionists, deny that the brain is the kind of formal symbol manipulator that classical cognitivists maintain it is.⁸ Instead, they suppose that, although the brain is a kind of computer, its computational processes are implemented by neural networks, and are, contrary to what classical cognitivists suppose, performed in parallel, at a sub-symbolic level.

My own concern, however, is with another kind of criticism. This criticism

⁵This view has been most prominently and systematically developed and defended by Jerry Fodor (Fodor; 1975), (Fodor; 1981), (Fodor; 1987). Steven Pinker is another prominent example (Pinker; 1997).

⁶(Churchland; 1989)

⁷(Ramsey; 2007)

⁸(Smolensky; 1988), (Churchland; 1989)

is about the role of meaning, or *representation*, in physicalist accounts of cognition. A lot of this has to do with what some have termed “The Problem of Causal Impotence”.⁹ This is the problem of showing how representational *content* – the purported meanings of the symbols – can be causally efficacious. Alternatively, one could say, it is the problem of showing how the familiar physicalist accounts of cognition can still leave room for content to serve a purpose in cognition.

As a starting point, we may draw attention to the implicit distinction, made in many theories of cognition, between representational *vehicle* and representational *content*. Simply put, it concerns the distinction between the symbol itself and what the symbol means. The problem of causal impotence arises from the fact that, on the standard accounts of computation, all the actual causal work is done solely by the vehicular properties of the symbols. In other words, what the symbols *mean* does, in this respect, not seem to have any relevance.

In response to this, some theorists have proposed that we really ought to eliminate the notion of representation in theorizing about cognition. One of the more early proponents of this view is Steven Stich who, at least at that time, proposed what he called a “syntactic theory of mind”.¹⁰ Briefly put, his idea is that we should just forget about the meaning of symbols entirely, and only focus on the actual system of syntactic operations that is implemented in the brain. Following Stich, many other theorists seem to have begun to entertain a non-representational understanding of cognition. This seems to have started, roughly, with Rodney Brooks’ article *Intelligence Without Representation*, and has most recently led to the formulation of a view called *Radical Enactivism*.¹¹

Ever since these non-representational views have come to the fore, there has been an ongoing debate about the role of representation in cognition.

⁹(Hutto; 1999)

¹⁰(Stich; 1983)

¹¹(Brooks; 1991), (Hutto; 2011), (Hutto and Myin; 2013), and (Degenaar and Myin; 2014).

Representationalists argue that representations are essential, whereas non-representationalists deny this. Much of the philosophical difficulties here, seem to have to do with the problem of how we might distinguish between systems that actually *do* employ representations from systems that don't. In fact, this may very well be impossible if, as some claim, the very notion of representation actually lacks a clear empirical operationalization.¹² Intuitively, I certainly see the difficulty of finding representation, as it were, “in the lab”. Perhaps, adopting a non-representational framework might be the appropriate thing to do after all.

Perhaps... But before we do so, I think it is important that we pay attention to a phenomenon, that most participants in the debate seem to neglect, consciously or unconsciously. A phenomenon that seems paradoxical in light of the idea that, as far as the the mechanics of cognition and behavior go, it is really only the vehicular side of representation that does all the important work (as some suggest). It concerns the fact that, subjectively speaking, *thoughts appear to have the property of being transparent*. What I mean by this, is that when we observe what our thoughts are actually like, “from the inside” as it were, it seems that what we find, is precisely the *content* of those thoughts, rather than any of the underlying vehicles that, supposedly, are “carrying” these contents.¹³ It is as if in thinking we are already beyond the thought, and with the thing itself.

Although I shall have to say more about this, intuitively, this seeming transparency is puzzling in light of the idea that the brain really is a “syntactically driven engine”¹⁴; an organ whose internal functioning proceeds on the basis of the local formal mechanics of its physiological parts. For it

¹²(Haselager et al.; 2003)

¹³But see (Churchland; 1985) for a dissenting view. The idea that mental states are transparent is usually mainly applied to *perceptual* states (Harman; 1990), (Tye; 2002). (This is often appealed to in order to show that we are never aware of any qualitative properties of experience, i.e. “qualia”.) But I think that the phenomenon of transparency carries over to mental states of a more “off line” variety as well. See also (Evans; 1982, p. 225) for what seems to be a similar view.

¹⁴(Dennett; 1998)

seems that in thought, we are actually directly aware of the *very referents* of thought. And thus it would seem that these very referents themselves, *non-local as they are*, have, in fact, a direct impact on us. For example, when I think about the moon, it seems as though the moon itself becomes somehow present to my mind. Or, at the very least, there is a very tangible feeling that, through the act of thinking, one actually establishes some sort of mental contact with the moon.¹⁵ (How else would you be able to think about *it*?) In effect, it appears as though one is directly aware of something that is spatially, as well as temporally, distanced from oneself. And thus, it appears as though one is directly *affected* by something that is itself non-local.

When we think about this transparent nature of our mental states, what readily comes to mind is the fact that our inner mental lives appear to be characterized by a certain sort of *openness*, that, indeed, is not, in any obvious way, suggested by the thought of formal mechanics driving computational processes inside the brain. Indeed, the mind appears to be *field-like* rather than point-like. It harbors a certain *wideness*, or extensiveness, that is difficult to reconcile with the image of 86 billion interacting neurons packed together in a small area of 1273.6 cm³. (Who, in his or her rational mind, would ever expect that a particular house, upon entering it, turns out to be much bigger on the inside than its outside appears to allow!)

What is, moreover, characteristic of this “mental field”, is that, through the act of thinking, it apparently manages to stretch itself far beyond the spatiotemporal confines of the bodies’ own physical embeddedness. Indeed,

¹⁵Dreams, I believe, also illustrate this phenomenon very well. For example, suppose that you have a dream in which you talk to a loved one that has already passed away, and that you deeply miss. Doesn’t it seem to you as if the person him or herself is actually there? In any case, most of us seem to highly *value* such dreams. This, it seems to me, has a lot to do with the fact that, in dreams, our experience is that they actually seem to be there in person. When we dream, it is not as though we are merely talking to mental images *of* the deceased. In that case, the significance and value of such dreams would probably not be that great. Moreover, I predict that even in case one becomes “lucid” during such a conversation, most of us would still treat the appearance in the dream as though it was the actual person him- or herself.

one could argue that thinking just *is* dissociating oneself from one's current physical surroundings. For in thinking we essentially seem to turn our mental regard to someplace and sometime else. This, at the very least, seems what the bulk of our conscious thinking processes are like.

Not only are we able to think about (or should we say “think *towards*”?) the past and the future, we also have the capacity to conceive of *alternative* or *hypothetical* realities. Moreover, through the act of thinking, we seem able to get into view, things that allegedly exist outside of space and time altogether, such as mathematical objects. Our mental capacities also seem to furnish us with a grasp of things of a specifically more *meta-physical* nature as well. For example, we seem to have some sense of the world as a whole, and seem able to contemplate such things as totality. Indeed, through our thoughtfulness, we even have the ability to cognize, and appreciate, the bare fact of reality, or existence, as such. So in various ways, thought literally appears to *broaden our awareness*, giving our sense of self and the world a more universal character.¹⁶

Thought, then, appears to have this strange, almost magical, ability to transcend, or dissociate itself, from the local physical situatedness of the body, from which, so we suppose, it takes root. And so, in so far as the things that we think about have a causal influence on our behavior, it would seem that the factors that actually shape our behavior “overflow” the local processes that constitute the brain, the body, and the immediate environment. In other words, the causal influence of thinking on the body seems to

¹⁶All of which, I believe, is the true significance of our folk-psychological conception of the mind. Contrary to “a mere animal”, whose consciousness and behavior is (supposedly) fully confined and determined by what is “immediately given” – it has a mere momentaneous consciousness – human beings are able, thanks to their “higher mental powers”, to seemingly transcend these limitations, and to enjoy a sense of themselves and the world that is more universal, and to have the freedom to consider alternative realities, and future possibilities, and to act according to them. Of course, whether this is all true, is a different matter. All in all I think it gives an intuitive picture of how we see ourselves and our place in the universe.

imply, essentially, a non-local form of interaction.

But, of course, these may be just appearances. And, perhaps, *they better be*. For granting them a reality within the physical world brings with it various difficulties. For one, it appears to require that our behavior is under the direct influence of non-local factors. That is, it suggests some sort of non-local interaction, or *action at a distance*. And this violates the principle of locality, so seemingly dear to our rational minds, according to which all that happens in the world reduces to contiguous physical interactions. One could also say that the previous story would suggest, perhaps to a somewhat naïve mind, that thoughts are actually instances of extra-sensory perception, or remote viewing. Ordinary thinking would turn out to be a paranormal phenomenon! So perhaps our mental lives cannot be what they appear to be after all: instances of disembodied mind-traveling across space and time.

I suspect that much of the worries that philosophers and other theorists have about representation stems from a conscious or unconscious awareness that, somehow, representation requires a non-local, or non-mechanical, sort of physical interaction. In any case, there seems to be an awareness that, for some reason or other, there is something problematic about the idea that a physical system employs representations. As William Ramsey puts it:

The difficulties associated with these efforts reveal just how tough it is to accommodate both our fundamental understanding of representation, and our desire to provide a naturalistic account of cognition. Given the central importance of representation throughout cognitive science, it is no small embarrassment that we are still unclear on how representations are supposed to operate, as such, in a biological or mechanical system. (Ramsey; 2007, p. 221)

The fact that so many theorists are enthusiastic about the prospect of giving an account of cognition in purely “form-sensitive” processes is, I believe,

indicative of this difficulty. For it seems that the only way to make intelligible how a physical system could operate with representational items, would be by showing how its representational activity reduces to that good old kind of mechanical, form-driven, interaction that we suppose to exist between physical items, and that we all feel to intuitively understand. However, as the problem of causal potency indicates, it seems that such a reduction basically *obliterates* the very *need* for mental content (why should the brain be *told* if it can be *forced*?). Indeed, this seems to be the prime reason why some, especially the radical enactivists, have come to propose, by way of a “can’t have” and “don’t need” policy with respect to representation, a non-representational account of cognition.¹⁷

I respect the skeptical attitudes that some theorists seem to have about “off line” mental imagery (which is, sort of, the kind of phenomenon that I have been referring to). And I understand that people may have certain reservations about my specific interpretation of it. As a matter of fact, I understand and agree that appealing to mental imagery is a delicate business.¹⁸ Perhaps this is why most non-representationalists are so remarkably hushed up about mental imagery. Or maybe it is because non-representationalists think that mental imagery is just a marginal phenomenon, or even an epiphenomenon, with respect to the actual accomplishments of cognition or intelligent behavior. And maybe, just maybe, a kind of dispositional account of the apparent transparency of thought, something along Dennettian lines, might turn out to be the most plausible after all, in which case it seems that we really don’t need to worry about representation, or the kind of weird interaction that representation seems to imply. In that case we can simply do away with the ontological and metaphysical problems of folk psychology, and safely say that human behavior is mechanistic through and through.

I am excited to inform that the problem is not that easily dispensed with. For it seems to me that, in place of mental imagery, there is another

¹⁷(Hutto and Myin; 2013, p. 22)

¹⁸See e.g. the arguments in (Dennett; 1991)

fact about our mental lives that, I think, is at least as problematic to a non-representational account of cognition as mental imagery appears to be. A phenomenon that is, in certain relevant respects, *rather like* mental imagery, but which is a lot more difficult to dispose of as illusory. Hence, much more difficult to “eliminate”. That phenomenon is *ordinary, plain, visual perception*.

Vision, it seems to me, *is essentially a representational process*. Indeed, the very *purpose* of vision is to *re-present*. To use a term of Andy Clark and Josefa Toribio, vision *stills*, for the organism, a certain “*representation hunger*”.¹⁹ In making this claim, I am inspired by William Ramsey’s book *Representation Reconsidered*.²⁰ One of the main points of that book, is that labeling something as ‘representational’ (as philosophers, neuroscientists, and cognitive scientists often do), requires a special sort of justification; it requires answering what Ramsey calls the “job description challenge”.²¹ That is, it requires that one shows that the thing that one labels as ‘representational’, actually has some sort of representational *function* or *purpose*. As Ramsey has convincingly shown, much of what is customarily labeled ‘representational’ is not actually deserving of that name. However, although many other things may fail to fit the bill, I believe that it can be easily shown that vision is, straightforwardly, a process of representation, and hence, deserving of the name ‘representational’.

Of course, in addition, the thesis that vision is a representational process comes with an important implication; namely that certain current non-representationalist accounts are inadequate (or at least *misleading*). Moreover, it appears to me that vision, partly due to its transparent nature, similarly appears to resist being given a mechanical interpretation. And this, I think, is significant given that non-representationalists are often motivated by “naturalistic” considerations. Of course, in what follows, I shall try to clarify all this a bit more. However, my main purpose shall be to show why

¹⁹(Clark and Toribio; 1994)

²⁰(Ramsey; 2007)

²¹(Ramsey; 2007, Ch. 1)

we should think of vision as essentially a process of representation.

II

The topic of this paper is visual perception; something that has always, in one way or another, fascinated philosophers. This is because there is much that is philosophically interesting, and perplexing, about visual perception, and perception in general. First – and I apologize for the fact that I am going to *digress* a bit – there is the bare fact of perception *as such*; the fact *that* perception occurs *at all*. For whatever else is true, it seems that, perception, as a natural phenomenon, has a deep metaphysical significance. A metaphysical significance that, I maintain, cannot simply be reduced to, or eliminated in favor of, something else (e.g., some idiosyncratic biological interest).

To appreciate why, we need to pause and reflect, first, on that ultimate riddle that is the truth: the fact that, rather than sheer nothingness, there is, instead, a strange, seemingly infinite, universe. (Which, of course, is the world that we live in.) And I cannot speak for others, but when I deeply contemplate this fact, there often arises a special feeling within me. A feeling of wonder that appears to have a universal and almost impersonal quality to it. As if the subject of that feeling was *existence itself*, rather than my own personal self. But whatever the true nature of that feeling (perhaps I am just delusional), I believe that we can all recognize that there is something really quite extraordinary about the fact that there is something, namely *the universe*, rather than nothing.

But let me not, here, dwell on this further. My intention in mentioning this conundrum of ontology is, in the first place, to say something about the significance of *perception*; the fact that there is perceiving at all.

To be sure, when it comes to the question of why there is something rather than nothing, the mere existence of a rock would seem as enigmatic as the existence of a perceptual experience. As Wittgenstein has said “Nicht *wie* die

Welt ist, ist das Mystische, sondern *daß* sie ist”.²² Yet despite this, I do think that, in this respect, there *is* something rather special about perception. For it is only when some sort of perceiving takes place that this solitary fact, this exception of exceptions, namely, the fact that there is something rather than nothing, actually becomes *apparent*. To borrow a term from another German philosopher, only when there is perceiving can there be some sort of “*Seinsverstehen*”.²³ Only through perceiving, or sentience, can the world make its presence felt. And one is inclined to say that such “understanding of being” (or the feeling of being) is in fact rather crucial. Perhaps even *essential* to the very mystery. For a world that goes altogether unnoticed, may in fact rather be closer to *nothing* than to something.²⁴

Thus, it is only through perception that the existence of the world itself becomes articulate. In a sense, one could say that perception is, quite literally, *the very significance of* existence. After all, in perception there literally is a “signaling of” existence – of existence and such, and of particular things that exist. It is only when some sort of perception takes place, that an actual registration can occur of the peculiar fact *that indeed* things exist at all. And since a world whose existence goes altogether unnoticed entirely might, for all that matters, not be at all, the bare fact of perception seems, in itself, something of special import.

But what is more, the fact that there is perception at all has, in recent years, come to acquire an *additional* philosophical significance as well. One that seems altogether different from what I have previously said. For under the assumption that perceptual states – “*phenomenal*” states – are merely derivative of more fundamental, non-perceptual “physical” processes, it has seemed to many that, in fact, there doesn’t seem to be a *sufficient reason*

²²(Wittgenstein; 1989, p. 150)

²³(Heidegger; 1967)

²⁴Really, try to imagine what it would be like for the universe to be, the stars and planets etc., without there being any entities to perceive it. Although *we* might be able to imagine such a world, for that world itself, there really is “nothing it is like” to be! Indeed, for all that matters to it, it might just as well *not* be at all!

why perceptual events are indeed happening within our physical world. I am of course referring to what David Chalmers has termed the ‘Hard Problem of Consciousness’.²⁵ It concerns the puzzle of why perceptual states should “arise from” or “be identical with” certain physical processes at all, and what their causal efficacy with respect to the physical consists in (if, indeed, they *have* any such efficacy).

So there is an interesting story to be told about the fact that there is perception at all rather than not. For only through perception is the existence of the world allowed to become apparent. And so if the fact that there is something at all rather than nothing can indeed rightfully be regarded as the ultimate mystery, it is thanks to perception that this mystery of mysteries can be *regarded at all*. But strangely enough, the fact there *is* indeed such a consciousness, appears itself, for seemingly different reasons, equally inexplicable! How something conscious could ever come from something non-conscious, seems as incomprehensible as something coming from nothing. This, at least, is how I like to think of it. So much for my digression.

I guess there are several reasons why I make all these seemingly unrelated comments. One reason, I suppose, is that I simply wanted to put some of my cards on the table. As can be inferred from the way in which I have framed some things, I think that thinking about perception really is thinking about “conscious experience” or “phenomenal consciousness”. In other words, I don’t think that there is much sense in thinking about perception as something that may occur *without* there actually, in some way or other, subjectively *appearing* anything, to something or someone. If there is to be genuine perception, it must somehow leave its phenomenal trace on the world. Moreover, I am convinced that we are still pretty much in the dark about how exactly consciousness might come about in a purely “physical” world.²⁶ Or how, more generally, something conscious intelligibly follows from some-

²⁵See e.g. (Chalmers; 1996), (Chalmers; 1998). See also the collected papers in (Shear; 1997).

²⁶Whatever “physical” really means (Stoljar; 2010).

thing non-conscious. I guess such mysterianism is not very respectable at the moment. Nevertheless, I do think that consciousness really is a very strange and puzzling phenomenon.

But what is more, I don't think that it is a *coincidence* that, not only we have difficulties with accommodating consciousness within our current "physicalist" worldview, but that the same kind of difficulties also appear to arise with respect to mental representation. I don't think it is accidental that both consciousness and representation appear to resist being given a proper physical interpretation. In fact, it seems to me that, in many respects, the problems of representation and the problems of consciousness are really the same. The difficulties that arise as soon as we try to get an account of just how they are to have their impact on the physical world, are good examples. Indeed, the problem of the causal potency of content is reminiscent of the issue whether or not "qualia" are epiphenomenal.²⁷ In any case, my intuition is that the problems of mental representation and consciousness are intimately linked.

This latter point is important, for it has become custom to separate the problem of naturalizing "meanings" from the problem of naturalizing "raw feels".²⁸ But to repeat, I believe that these two facets of our mental lives are in fact deeply intertwined. I sincerely doubt that we can separate the problems of representation from the problems of consciousness.²⁹ This is

²⁷See e.g. (Jackson; 1982), and (Chalmers; 1996, Ch. 4 & 5).

²⁸See e.g. (Fodor; 1987, p. 16-18) and (Güzeldere; 1997, p. 22-24).

²⁹See e.g. (Horgan and Tienson; 2002) and (Mendelovici and Bourget; 2014) for similar views. I guess John Campbell's *Reference and Consciousness*, in which the thesis is put forth that the very capacity for reference is essentially grounded in conscious experience, also provides a good example of this line of thinking (Campbell; 2002). As I come to think about it, I might as well put a reference to (Harnad; 1990) and (Barsalou; 1999) here as well. Although, in their writings, these authors are not particularly concerned with the metaphysical relationship between the phenomenal and the representational (not explicitly at least), they do seem to recognize the fundamental importance of perceptual content to cognition. Note, however, that these are all views that mainly aim to ground *representation* in *consciousness*. However, there are also theorists that have argued that it is in

indicated by the fact that, as soon as we try to apply the notion of mental representation *outside* the domain of conscious-level experience, and try to apply it to more straightforwardly physical things and processes, we seem to run into all sorts of difficulties. At least, this seems to be what many of the recent discussions of mental representation have now made clear. I shall say more about this shortly.

First of all, we must note that *the very notion* of mental representation is, first and foremost, derived from our own conscious experience. As we become aware of our own mental life, from the first-personal point of view, and become aware of how conscious thought takes place, we come to discover that, apparently, we have an in-built ability to mentally represent things to ourselves. As Angela Mendelovici and David Bourget note, representation “...is a phenomenon that we notice in ourselves at the first instance...while states that we cannot introspect upon might have intentionality, our reference-fixing grip on our target is through introspection” (Mendelovici and Bourget; 2014).³⁰ So although it may now be common, especially in the cognitive sciences, to attribute representational properties to unconscious sub-personal processes, we know representation, first and foremost, as a phenomenon of conscious experience.

fact *consciousness* that is to be grounded in *representation* (Dretske; 1995), (Tye; 1995). According to this view, conscious experience is basically a specific form of representational processing. Typically, it is committed to the idea that there is such a thing as unconscious representation. Moreover, from what I reckon, proponents seem to be largely motivated by the idea that, as computers show, unconscious representation is a fairly non-mysterious phenomenon. And this seems to invite optimism about the prospect of giving an equally non-mysterious explanation of consciousness. I shall have to say more about this view later.

³⁰Although Mendelovici and Bourget are writing about “intentionality”, it can be inferred from the article that we can safely substitute here ‘representation’. As a matter of fact, the majority of authors in the field simply equate the concept of intentionality with the concept of mental representation; to be in an intentional state just *is* being in a state of representing some thing or state of affairs, and *vice versa*. (See also footnote 6 of (Colombo; 2014).) Mendelovici and Bourget’s story about reference-fixing can safely be understood as a story about the initial reference-fixation of ‘mental representation’.

Of course, from this it does not automatically follow that there can be no mental representation without consciousness; that the concept cannot find some legitimate application outside the domain of conscious experience. As I have just pointed out, it is very common in cognitive science to posit the existence of sub-personal mental representations inside the brain that are not accessible for conscious inspection. However, it seems that the more theorists have begun to think about what this could possibly mean, the less credible this view has become. As Daniel Hutto writes:

At root, there must be some genuine way to explain what informational ‘content’ is and what it is for subpersonal mechanisms to ‘communicate’ with one another in terms of it... Otherwise it seems we must accept that talk of the reception, manipulation and production of content at this level is, as McDowell (1994/1998) noted some time ago, “irreducibly metaphorical” (p. 349). (Hutto; 2013)³¹

One of the main difficulties here, it seems, is that when we look at the kind of sub-personal processing that the brain is engaged in, all we seem to find are just ordinary causal mechanisms. At the physiological (or computational) level, we see that our senses are stimulated in all sorts of complex ways, and we see that, in response, the brain engages in all sorts of intricate neurological processes. But these are fairly ordinary physical interactions. However, in order for representation to be sufficiently real, it must actually have some proper causal role within the system. And the intuition is that, in order for

³¹To be sure, Hutto is referring here to (McDowell; 1994)/(McDowell; 1998). As a matter of fact, I think it is worthwhile to let McDowell speak for himself: “The “sub-personal” account of a sensory system, which treats it as an information processing device that transmits its informational results to something else inside an animal, cannot adequately characterize what its sensory systems are for the animal (as opposed to what they are, *metaphorically speaking*, for the internal parts that receive the results of the information-processing): namely, modes of sensitivity or *openness* to features of the environment” (McDowell; 1994, p. 197) (my emphasis). It appears that my views and McDowell’s are, in this respect, rather similar.

representation to have such a causal role, something more is needed than just ordinary mechanism. As Haselager *et al.* note:

Importantly, there must be something more than the mere operation of causal forces in order to warrant talk of representations. That is, the standing-in, the content, must play a crucial role in the system's functioning. But specifying exactly what this crucial role of the content is, and how it can be discovered that the behavior of the system relies on this crucial role of content when examining a system's causal-physical operations, turns out to be surprisingly difficult. (Haselager *et al.*; 2003, p. 19)

It is true that sometimes specific parts of the brain respond in highly exclusive ways to only very particular environmental stimuli. And it is common to suppose that, therefore, when the brain responds in this way, the brain is in a representational state of representing that particular environmental stimulus. However, as has been pointed out, such highly specific correlations are not sufficient for there to be instances of representation.³² In fact, it seems enough to suppose that, at this level, all there is are simply highly complex physiological processes that causally mediate, through sheer mechanical force, the organism's responses to the environment. There doesn't seem to be any functional merit in supposing that any of the components that make up these various causal chains are also representing anything.³³ Again, why, at this level, should the brain be "told" or "conveyed" anything, if all its proceedings already happen through sheer mechanical force anyway? In order for representation to take place, it seems that something more is needed. But it is difficult to see what, physiologically speaking, this extra could be.

To be sure, William Ramsey has somewhat convinced me that there *might*, in fact, be a genuine and non-metaphorical way of thinking about sub-

³²See e.g. (Ramsey; 2007, Ch. 4), and (Hutto and Myin; 2013, Ch. 4).

³³See e.g (Van Gelder; 1995), (Chemero; 2000), and (Chemero; 2009).

personal (i.e., non-conscious) representation after all.³⁴ Consider a pocket calculator: a purely mechanical device that is apparently able to be engaged in intricate mathematical problem solving. According to Ramsey, in order to make even *intelligible* how a pocket calculator manages to solve the kind of problems that it does, one must suppose that things that are internal to the pocket calculator actually function as a kind symbols. But if it is OK to attribute symbolic properties to wholly non-conscious things inside a pocket calculator, then surely a similar thing might perhaps be possible for sub-personal processes in the brain. Especially so if this gives us a physically acceptable account of how the brain manages to engage in certain types of problem-solving.

In this respect, I have also learned from Tim Crane.³⁵ We possess a lot of propositional knowledge about the world, but we are not all the time consciously entertaining these propositions in our minds. According to Crane, it is plausible to suppose that this propositional knowledge is nevertheless somewhere present within us even though we are not consciously aware of it. So the notion of an unconscious representational item or process appears not that crazy after all. Also, we might add that, very often, solutions to problems simply *come to us* in the form of a conscious thought, making it appear that the brain was all the time, much like a pocket calculator, sub-consciously working on the solution – the conscious thought that is finally “displayed” in our consciousness. (Sometimes we notice that this process falters: when for example we feel a word on the tip of our tongue, but it nevertheless doesn’t manage to get through to us). So perhaps representation really are sub-consciously present in the brain in the same way as they are in a pocket calculator.

Let me try, as briefly as I can, point out my reservations with respect to this kind of thinking. First of all, Ramsey’s story essentially depends on the idea that the kind of mathematical problem-solving that a pocket calculator is

³⁴See especially (Ramsey; 2007, Ch. 3).

³⁵(Crane; 1995, Ch. 1)

engaged in is an *intrinsic* property of the pocket calculator (Ramsey; 2007, p. 96-102). But I think that this is a misconception: it is only because the pocket calculator is *used by us* in a certain way, that its intricate workings become instances of calculation.³⁶ So I would say that the symbolic properties and powers of the pocket calculator are ultimately derived from the fact that it is being used in a certain way by already conscious minds. Therefore, a pocket calculator is an inappropriate model for understanding how the brain could host and operate symbolic items at the sub-personal level (unless we want to say that their status as symbols is also, in a sort of top-down fashion, derived from a conscious mind). Moreover, as Ramsey himself acknowledges, the idea that the cognitive architecture of the brain is like that of a digital computer is losing support anyway.

This brings me to Crane's point. Very briefly: I agree that there must be some mechanism for propositional memory and the way in which conscious thoughts coherently and functionally enter into our conscious awareness. However, why should we suppose that the brain needs unconscious thoughts to algorithmically "cook up" an appropriate conscious thought? Indeed, connectionists models may explain how the brain does this without using prior representations.³⁷ I might add to this that representation is what representation does. And so perhaps only when they are consciously processed can representations do their representational work. At best, we might say that the kind of things that are supposed to lay dormant within us, and that explain how propositional knowledge pops-up into conscious awareness, are instances of *proto-representation*, rather than instances of actual representation.

In any case, it seems to become increasingly clear that the notion of sub-personal representation (in the materialist-computational sense) has all sorts of problems. Not only does it seem difficult to come up with an account of how exactly the designated processes can be operative *as* representations, the

³⁶Unfortunately, I cannot afford to give my own arguments for this here, but see (Horst; 1999).

³⁷An idea that seems to be suggested in (Ramsey; 2007).

question is also whether there is even any explanatory utility in supposing that instances of representations *are* operative in an organism. However, if these difficulties cause one to be a total skeptic about mental representation, I believe that one is simply looking in *the wrong place*. My advice: *come to your senses!* For I take it that, at the level of conscious (visual) experience, it is, when properly understood, pretty obvious that there is going on a process of representation – *how* this can be, is of course, a different matter. (Also, I would say that for explanatory purposes, the supposition that, in vision, a process of representation takes place, is also in fact quite essential.³⁸) Fortunately, apart from a few hard-headed skeptics, the majority of theorists already seem to agree with this.

Furthermore, as I have already pointed out, the very concept of mental representation is derived, in the first instance, from our own conscious experience anyway. So perhaps the various attempts to try to look for it somewhere else were misguided all along. In any case, whether there actually is or is not some kind of unconscious, sub-personal, representational processing, we must, at some point or other, going to take stock of mental representation as a phenomenon of conscious experience. As far as I can see, that is where the real magic happens.

This gives me an opportunity to say something about the way my own thesis should be appreciated within the larger debate. For it seems that, in the literature on representation, there are in fact two closely related, yet distinct, issues. And failing to keep them distinct may result in ambiguity with respect to what exactly it is that is being claimed. Therefore, I think it is important to put some emphasis on it.

First, there is the issue of whether there is such a thing as mental representation *at all*. That is, it concerns whether or not with respect to the *overall* cognitive and behavioral aspects of an organism, some representational process is taking place. But this should be contrasted with another issue that theorists seem to have been concerned about. This is the issue

³⁸I shall say more about this in the final section of this paper.

of whether these more “global” processes of representation (if, indeed, there are such) should be *explained* in terms of *other, more primitive* processes of representation.

To illustrate, as William Ramsey points out, it is generally agreed that cognition is best described “as the conversion of representational inputs into representational outputs” (Ramsey; 2007, p. 69). So much is, according to Ramsey, uncontroversial. However, what *is*, according to him, controversial is whether this process of conversion also itself breaks down into smaller, intermediate, representation conversions. An example of a view that affirms the latter, is David Marr’s computational theory of vision.³⁹ According to Marr’s theory, the conscious visual percept, which according to this theory is also itself representational, is in fact constructed out of more primitive, sub-personal, representations.

The idea here is that underneath the surface of conscious-level experience, there is all kinds of representational activity going on, from which conscious-level representation is merely the end product. However, it is surely possible to maintain that vision *as a whole* is a representational process, but deny that this process itself is also in fact *accomplished* by other representational processes. Doing so leads to a view that some recently have referred to as “non-inferential representationalism”.⁴⁰ In other words, one can maintain that genuine representation occurs at the level of conscious experience, without maintaining that there is representation at the sub-personal level. To conclude, the following comment captures perhaps somewhat more the general spirit of what I have in mind:

³⁹See e.g. (Poggio; 1981) and (Shapiro; 2011, Ch. 1) for a summary of Marr’s theory.

⁴⁰See (Orlandi; 2014, p. 35). The radical enactivists seem to refer to this view as “maximally minimal intellectualism” (Hutto and Myin; 2013, ch. 5-6). It should be noted that these “minimal” representational accounts of mentality, including Orlandi’s, basically rely on the idea that phenomenal experience is intrinsically “contentful”. However, as I shall argue later, the representational nature of vision is quite independent of whether or not phenomenal vision is in and of itself “contentful”. So one could go even more “minimal”, and argue that although vision is not “contentful”, it is nevertheless representational. I think this poses a serious problem for the radical enactivists. But more on this later.

So, even if in some cases one could view a system *as a whole* as representing something, it does not follow that one needs to posit “stand ins” in the causal explanation of that particular behavior. If it can be shown that the system’s state is best seen as an *emergent property of the overall operation of the system*, a non-representational account seems the most natural way of analyzing the system. (van Rooij et al.; 2002, p. 351) (emphasis from the original text)

Although I do not agree with everything that is being said here (I think that representation *is* causally relevant in producing behavior), there nevertheless is a lot in this statement that I find appealing. First of all, there seems to be in it a recognition of the idea that mental representation will not be found in any of the physical sub-parts of an organism. Rather, it is only with respect to the organism *as a whole* that it makes sense to talk about representation. This accords well with my own view that mental representation happens, first and foremost, at the level of conscious experience. For conscious experience is also something that essentially applies to the organism as whole, at the level where the organism is an actual subject or individual.⁴¹

But what is more, the previous comment speaks of representation as an *emergent property*. To be sure, I am not quite sure whether the notion of emergence is actually a workable concept.⁴² But if, as it seems, mental representation suggests a special kind of interaction, something different from (or irreducible to) the kind of contiguous micro-physical interactions that we are familiar with (as indeed I have hinted at in the previous section, and which also seems to be recognized by others), then the idea that perhaps representation is some kind of emergent property with special causal powers,

⁴¹To be sure, by “subject” I don’t necessarily mean something that is essentially merely passively undergoing things. Rather, by “subject” I simply mean something of which certain indexical statements are true.

⁴²See for example (Kim; 2006) for a critical assessment.

naturally suggests itself. This is especially so, given the supposition that mental representation is a feature of conscious experience. For consciousness has also been thought of as an emergent phenomenon.⁴³

But let me go into this later (in the final section of this paper). For now, what I want to make clear is that my thesis only concerns vision as a whole; the process of vision in its entirety. In other words, my claim is that vision, *when taken as a whole*, is essentially a process of representation. And in my mind, this basically means that we should locate the representational aspects of vision primarily at the level of first-person conscious experience. My claim *does not* include the thesis that vision can also be explained *in terms of further* sub-personal processes of representation.

Now, if my claim is merely that vision *as a whole* is a process of representation, whereas the really *controversial* claim is, as Ramsey seems to maintain, the claim that the totality of vision itself, in fact, further breaks down into other, more low-level, representational processes, then it might appear that what I am proposing is really not all that interesting. For it would seem that what I am claiming is merely what everyone else already thinks anyway.

Yet it seems to me that there certainly are theorists, most notably the radical enactivists, that *do* in fact vehemently reject the idea that perception, vision, is a process of representation.⁴⁴ In other words, they don't merely reject the idea that vision is the result of sub-personal transformations of sub-personal representations. They also reject the idea that vision *as such* is a representational process. Indeed, they seem to be rather serious about this.

In addition, what I *mean* when I claim that vision is a representational process, may, in certain respects, be somewhat different from what is usually meant by this. A similar thing goes for the *reasons* that I have for making this claim. For example, the fact that vision is representational, in my sense,

⁴³See e.g. (Chalmers; 2006).

⁴⁴(Hutto and Myin; 2013, Ch. 5-6)

is quite compatible with the thesis that perception is “direct”, and the thesis that phenomenal experience is not inherently “contentful”.⁴⁵ Indeed, the reasons for motivating that vision is a process of representation are much more straightforward. So I do think I have something new and interesting to bring on the table here.

Some will object that, by changing the meaning of theoretical concepts in this way, it is easy to make a point and win an argument.⁴⁶ However, I think this criticism is only valid when the meaning of an already familiar concept is changed beyond recognition. Yet I believe that my own use of the concept of representation is, as a matter of fact, pretty much in accordance with the way in which, pre-theoretically, we understand its meaning. And as William Ramsey has pointed out, when it comes to the issue of whether or not one makes legitimate use of the notion of representation, this is what counts.⁴⁷

But what is more, I believe that my account of the representational nature of vision also respects very much our intuitions about what it is about representation that makes it so problematic from a naturalistic point of view. That is, my account makes clear that, even if vision is “direct” and “contentless”, it nevertheless behaves in the exact same problematic way that you would expect instances of representation to behave. Now, I think that *independent of that*, there are also *good reasons* to refer to the whole process of vision as ‘representational’. But I guess people may differ with me on this point. The terminological side of the whole issue here is, I believe, not what ultimately matters. For whether or not one agrees that there is an appropriate sense in which vision is deserving of the name ‘representational’, the result is the same. Even if, according to the standard definitions, vision ought not to be considered representational, we are still left with essentially

⁴⁵What is *usually* meant by the proposition that vision is representational, is either that visual perceiving is “indirect” or that visual perceiving is intrinsically “contentful”. However, I shall argue that there is an appropriate sense in which vision is representational, even if vision is “direct” or “contentless”. But more on this shortly.

⁴⁶(Chemero; 2009, p. 66)

⁴⁷(Ramsey; 2007, Ch. 1)

the same kind of problems that we would be left with if vision *were*, in this more traditional sense, representational. Nothing is solved by denying that vision is “indirect” or by denying that vision is intrinsically “contentful”.

III

The basic idea that I am going to put forward is fairly simple. As I have already said, my thesis is that vision is essentially a representational process. Visual perception serves a *representational function* within the cognitive and behavioral economy of the (human) organism. Again, to use the words of Clark and Toribio, vision basically stills, for the organism, a certain “representation-hunger”. But let me now also tell *why* I think this is so: vision presents the organism, *facilitates the presentation of*, things and/or events that are, in a relevant sense, still *absent* (or *non-present*). *Vision makes present, for the organism, what would otherwise still be absent*. One could say that sight really is *foresight*. Or that viewing is, in a way, really *remote viewing*. For vision essentially provides a glimpse into what is happening *somewhere else*, or what is likely *yet to come*. In this way, vision endows the organism with extra *anticipatory power*. For reasons that, hopefully, will become clear, this allows one to say that vision is essentially a process of representation. And what is perhaps more important; all this is quite independent from whether or not vision is representational in the more traditional senses of that term. So let me go into these first.

When one reads the literature, one finds that the proposition that vision is representational can in fact mean different things. I have already distinguished the idea that vision as a whole involves representation, from the idea that the implementational mechanisms that *realize* vision are representational. Here, we are only concerned with the first of these two ideas. Our concern is with the idea that *actual visual perceiving* involve representation. But this also can mean different things. In fact, there commonly appear to

be, in the literature, two different views associated with the proposition that visual experience is representational. It will be worthwhile to say something about these two views. Not only because these two views often seem to be conflated, but because it will also make it easier to show that, in fact, there is a *another* more *direct* way to think about vision as a representational phenomenon.

First of all, the proposition that vision is representational, can denote the idea that, in visual experience, we are not presented with the external world itself, but rather with *something else* – presumably something wholly mind-dependent – that is a *representation of* the external world.⁴⁸ According to this view, our visual awareness of the external world is *indirect* and *mediated by* what we are presented with in visual experience. In other words, it supposes that there is a kind of “veil of perception” between subjects and the world. On this account, one could say, our perceptual condition is rather like seeing the world always only indirectly through a television screen (or if you will, via the mental screen of a Cartesian Theater).

The view that we are only “indirectly aware” of the “external world” is a relatively old and familiar idea in philosophy. It has often been referred to as ‘indirect realism’ or as the ‘representative theory of perception’. The philosopher Ned Block has simply referred to it as ‘representationalism’. However, he distinguishes it from a more recent view, which he calls ‘representationism’, which is the second view that I want to elaborate on.⁴⁹ As it happens, however, this view is also often, though somewhat misleadingly, referred to simply as ‘representationalism’. Misleadingly, because this view appears to be, as a matter of fact, *radically opposed* to the idea that we only perceive the world indirectly. Indeed, some even refer to it as “the modern day’s direct realism about perception”.⁵⁰

⁴⁸I suppose a notable example of this view can be found in (Jackson; 1977). More generally, sense-data theories of perception seem to fit the bill here.

⁴⁹See (Block; 2006). Influential formulations of representationism include, (Dretske; 1995) and (Tye; 1995).

⁵⁰See (Aydede; 2014). Michael Tye, a prominent defender of representationalism, has

From the point of view of representationism, the claim that vision is representational basically means that in visual experience, the world is represented *as being a certain way*. That is, it takes visual experience to be inherently *contentful*. In a sense, it supposes that visual phenomenology is actually akin to a kind of *language*, to the effect that visual experiences actually “tell” the subject *that*, as a matter of fact, such and such is the case. However, the difference here (I would say) is that, on this view, it does so in a communicative medium that is essentially non-symbolical, in the sense that it presents the content of the message *directly*, rather than via a symbolic medium. This view also comes with the idea that experiences have conditions of satisfaction that are actually *built into* the experience itself. And so the idea here is that, *independent of how we might interpret or “take” them*, phenomenal experiences have, in and of themselves, a certain fixed meaning or “content”; a meaning or content that determines under which conditions the experience is satisfied.

To better appreciate what is being said here, it is worthwhile to consider how representationist accounts of hallucination typically differ from non-representationist accounts; accounts that *deny* that visual experience is inherently contentful.⁵¹ First of all, both views agree that cases of hallucination involve some sort of *error*. Both agree that in hallucination something that is false is erroneously taken to be true. What representationists and non-representationists disagree about is the *source* of the error.

According to representationists, when we hallucinate, the source or the error lies in the experience itself, or in the phenomenal appearance. That is,

also explicitly denied that there is any such thing as a veil of perception (Tye; 2014). Strangely enough, most philosophers don’t seem to be aware of the confusion that results in calling this view representationalism. It is comforting to see that at least Block, as well as some others, are aware of it. See (Robinson; 2008), (Wright; 2008), (Thompson; 2008), (Aydede; 2014), and (Lehar; 2015).

⁵¹See (Travis; 2004) and (Brewer; 2006) for two prominent examples. The radical enactivists also deny that experience is intrinsically contentful (Hutto and Myin; 2013, Ch. 5 & 6).

the phenomenal features of an experience are themselves directly assessable for truth or accuracy. One could say that, according to this way of thinking, each time something phenomenally appears in one's consciousness, some kind of *claim* is being made about the world. A special claim, however, in that it has a certain *autonomy* relative to our personal-level judgments. In other words, even if we know on an intellectual level that the pink elephant that we seem to see is not really there, our perceptual systems may still continue to proclaim the contrary.

Non-representationists, however, typically claim that the source of the error is further "upstream", at the level where we make judgments and form beliefs *about* our experiences. According to such accounts, experiences simply are what they are and, therefore, not assessable for truth or accuracy. When we falsely believe that what visually appears to us is a pink elephant, then, on this account, it is ultimately *we* who are to blame, and not our visual apparatus. At most, our visual apparatus might *mislead us* into thinking, falsely, that something is the case. But again, *we* are to blame; the error results because *we* have wrongly interpreted things. The eyes themselves are, strictly speaking, "dumb". In and of themselves, they don't have any capacity to tell either truths or lies.

But let's forget about non-representationism, and focus on the difference between *representationism* – the view that phenomenal experience is contentful – and *representationalism* – the view that we perceive the world only indirectly. For the difference between them is rather subtle. As a matter of fact, I don't think it is even entirely inadequate to say that, in many contemporary writings, these two views are often simply treated as one. Indeed, one sometimes gets the impression that many theorists are not even aware that there *is*, in fact, a difference.⁵² No doubt this is largely due to the fact that *representationism* is often sold under the banner of *representationalism*, which, I suppose, most theorists naturally associate with the good old representative

⁵²A notable exception is Christopher Hill (Hill; 2009, Ch. 2). He gives, I think, a very good account of the conceptual differences.

theory of perception. However, there *is* a difference. A difference that might perhaps be explained in terms of the distinction between *content-based* and *vehicle-based* theories of mental representation.⁵³

According to vehicle-based representationalism, what we are presented with in experience are actually representational *vehicles*. According to this way of thinking, we could think of what appears to us as the phenomenal analogue of written expressions; a sort of wholly syntactic configurations of “phenomenal ink” (or “mental paint”, as Ned Block calls it).⁵⁴ On this view, what appears to us are mere *symbolic signs*. On the other hand, content-based representationalism is the view that, in experience, we are actually directly acquainted with the very *content* of the representation. Indeed, it supposes that in experience, the representational vehicles themselves are in fact *transparent* (an idea that I have already elaborated on earlier). On this view, experience is intrinsically language-like, though *non-symbolic* in the sense that it conveys its contents *directly*.⁵⁵

As far as I understand, one might say that good-old *representationalism* is essentially vehicle-based, whereas contemporary *representationism* is content-based. As it happens, both views have very different metaphysical and ontological implications. Let me conclude by mentioning some of them.

First, on a vehicle-based account, phenomenal character is *not* in and of itself “*world revealing*”. It doesn’t necessarily show us, in and of itself, what the world is actually like, qualitatively speaking. Because it conceives of phenomenal appearances in vehicular, and thus symbolic terms, it leaves open the precise nature of the relationship between what experience is like, and what the world is like. In fact, even if the world appears to us as it actually is, this would be, on this view, merely accidental. An analogy: although a painting may resemble the thing or situation that it represents, it is not *essential* for a painting to resemble what it represents. One could make

⁵³I borrow these terms from Brad Thompson (Thompson; 2008).

⁵⁴See (Block; 2003).

⁵⁵Some would say that it is also a language that is essentially *non-conceptual*. See e.g. (Tye; 1995).

an abstract or non-figurative work of art that represents the *exact* same thing or situation. Thus, from a vehicle-based perspective, the relation between the phenomenal character of experience and the character of the world, is just as loose as the relation between the character of a symbol and that which the symbol symbolizes.⁵⁶

Not so with content-based accounts. For on such accounts, phenomenal character *is* essentially world-revealing. On this view, phenomenal characters *directly specify* what the world is actually like. Moreover, it *denies* that experience is picture-like, in the sense that, for it, the “qualitative correspondence” that exists between the phenomenal character of experience and the world (at least in the veridical cases) is not, in the first instance, grounded in resemblance. Instead, phenomenal character is supposed to represent *intrinsically* or *directly*. For note that, whether or not something resembles something else is not an intrinsic property of that thing. At the very least, it seems that there is always some sense in which it is appropriate to say that everything resembles everything else, whereas there is also always a certain sense in which it is quite appropriate to say that, in fact, nothing ever resembles anything else. According to a content-based interpretation of experience, however, the nature of phenomenal experience is such, that it represents in a way that ambiguity, or indeterminacy, does not enter. Again, the idea is that experience represents directly; it supposes that experience represents in a *content*-based manner.

Finally, content-based and vehicle-based accounts differ with respect to their ontological implications. From a vehicle-based perspective, what appears must, in some sense, actually exist. A principle which has come to

⁵⁶To be sure, we might point out that, even if we suppose that phenomenal character is not in and of itself world-revealing, and suppose that all that features in experiences are various phenomeno-syntactic configurations, we might still allow that there can nevertheless (and, for functional reasons, probably *must*) be a kind of *structural isomorphism* between the phenomenal structure of experience and the structure of the world, especially if there are many fine-grained correlations between the two.

be known as the ‘phenomenal principle’.⁵⁷ However, from within a content-based perspective, this is usually denied. In effect, representationalists and representationists both give very different accounts of what happens in cases of illusion and hallucination.

To illustrate, from a vehicle-based point of view, our experience is rather like reading the newspaper. Although it may be full of lies, it is still the case that the symbolic expressions themselves really exist in the way they appear. (At least, from the fact that the newspaper is full of lies, it doesn’t follow that the newspaper itself, with its various claims, does not, *as such*, exist.) This is different with a content-based account. On a content-based account, experiencing is more like *believing*, in the sense that, just as one can believe that *p*, one can *experience* that *p*. However, just as it doesn’t follow from the fact that one believes that *p* that *p* is also *true*, neither does experiencing that *p* imply *p*. Thus, the fact that one experiences a red color somewhere in one’s visual field does not imply that there actually *is* a red color.⁵⁸

These, then, are the two most common views associated with the idea that vision is a representational process. To be sure, I have reconstructed these views in such a way that they are maximally distinct. However, there certainly seem to be mixed versions out there – i.e., accounts that suppose that the phenomenal character of experience is both characterized by vehicular *and* contentful aspects.⁵⁹ Moreover, it may very well be that representationalists and representationists disagree with the way in which I have reconstructed their views. Yet I think that, given our conceptual space, the content-based/vehicle-based distinction is a neat tool to distinguish between two very different ways to apply the concept of representation to vision. Also, dividing theoretical space in this way, is the only way I can think of to make sense of the claim that contemporary representationism à la Tye and Dretske

⁵⁷See (Robinson; 1994), and (Fish; 2010).

⁵⁸Whether this also makes sense is, of course, a different matter. See (Robinson; 2008) and (Thompson; 2008) for critical assessments.

⁵⁹I count as examples (Peacocke; 1983), (Block; 2003), and (Shoemaker; 2003).

is not just a version of the traditional representative theory of perception. (To be honest, at times, I am still at a loss as to what exactly the view is supposed to amount to.) But what is more, my point is that vision can be regarded as worthy of the name ‘representational’, even though it may neither be indirect nor intrinsically contentful. Indeed, it seems to me that vision’s status as a representational process is quite independent of this.

However, in order to appreciate this, I think it is important that we approach vision *phenomenologically*, rather than that we make the mistake of focusing exclusively on the “underlying” physiological mechanisms of vision. Of course, this has everything to do with my conviction that the representational features of vision are located at the level of conscious experience. In other words, we must consider what visual perceiving is actually *like*. This is not something that is easily revealed simply by looking at what goes on at the physiological level of description. Indeed, the picture that the physiological story suggests seems, as a matter of fact, even somewhat *opposed* to what is suggested to us phenomenologically. (Which seems part of the reason why, if we only look at it physiologically, we tend to become anti-representationalists.) But let me come back to this later.

Consider what it is actually like to see. An answer that readily springs to mind is that vision characteristically consists in the perception of *shape* and *color*. However, it seems to me that there is another important characteristic of vision; a feature that is deeply puzzling, and yet often paid too little attention to in philosophical discussions; it concerns the fact that, in vision, we are presented with objects and states of affairs that have their being *a certain distance away from us*. It concerns the fact that there appears to be a *spatial gap* between the seer and what it seen.

Stretch out your arm, and look at your hand. Does not your hand make its appearance, and does it not seem as though it does so *from a location* that is *a certain distance away from you*? In other words, although one’s *seeing of* the hand seems to happen “*over here*”, the hand that is seen, *the object seen*, nevertheless makes its appearance “*over there*”. Yet strangely enough,

“over here” seems to be directly aware of “over there”. And this seems to happen without “over here” needing to go “over there”, and without “over there” needing to come “over here” (after all, wouldn’t we otherwise see much more details?). In other words, one appears to see the distant object directly. However, in doing so, we ourselves and the object nevertheless *maintain our distance*. Neither the seer, nor the seen, crosses the gap that exists in between (hence the gap). And so indeed, given the phenomenology of vision, one is inclined to say that visual perceiving must involve some sort of “spooky” action at a distance.

Now if we take this phenomenon seriously, as I think we should, then I think that *already in and of itself*, it provides sufficient reason to regard vision as a process of representation. *Even if visual perception is neither indirect nor contentful*. In fact, in what follows I shall assume that vision actually is direct and intrinsically contentless.

What does it mean for something to represent? Really, what *is* representation? To be sure, there might not be any one set of sufficient and necessary conditions that define the nature of representation.⁶⁰ Indeed, according to some, representation really is a “cluster concept”.⁶¹ However, this does not mean that whether or not something falls under this cluster is an empty question. It *does* mean that whether something is an instance of representation depends on one’s concept of representation. But although there might be more than one legitimate way of using the term ‘representation’, it does not follow that, in this respect, anything goes. As William Ramsey has convincingly shown, there certainly are proper and improper ways of using the concept of representation.⁶²

As I have already briefly mentioned, in the neurosciences and in cognitive science, it is common to regard certain patterns of brain activity as forms of mental representation, because there exists a systematic correlation

⁶⁰(Stich; 1992)

⁶¹(Ramsey; 2007, Ch. 1.1)

⁶²(Ramsey; 2007)

between these patterns of activity and certain features of the environment. The common assumption is that such patterns *represent* these environmental features. But as others have pointed out, the fact that a specific form of neuronal activity takes place exclusively in response to some particular environmental stimulant, is not *sufficient* (and, it would seem, neither necessary) for that neuronal activity to represent that environmental stimulant.⁶³ In other words, the fact that the presence of a pink elephant elicits a specific type of neuronal activity in V1 does not, in and of itself, suffice for that pattern to represent a pink elephant. Something more substantial is required for something to be an instance of representation.

But what? Again, Ramsey has useful things to say. According to him, there are basically two ways in which such a neuronal pattern of activity can in fact turn out to be an instance of representation. One of these I have already briefly mentioned, and that is that the neuronal pattern in fact instantiates (part of) a process of formal symbol manipulation. The second way is for the neuronal pattern of activity to have a certain *structural resemblance* to the thing or state of affairs that it represents, a structural resemblance that must, moreover, *as such* be used by the cognitive system, and which must also be used *because* it bears this resemblance to the thing or state of affairs in question. Ramsey refers to these as “io-representation” and “s-representation” respectively.⁶⁴

Although Ramsey has his reasons for distinguishing between these two forms of representation, in my mind, they essentially come down to the same principle.⁶⁵ This is because the very *reason* that processes of formal symbol manipulation work, is precisely because the symbols, and the system of syntactic operations that may be performed on them, bear a certain structural

⁶³See (Ramsey; 2007, Ch. 4) for an extensive treatment. The radical enactivists refer to this as the “Covariation Doesn’t Constitute Content Principle” (Hutto and Myin; 2013, p. 67).

⁶⁴See (Ramsey; 2007, Ch. 3) – “io” stands for ‘input-output’ and “s” stands for ‘simulation’.

⁶⁵(Ramsey; 2007, p. 102-104)

resemblance or “isomorphy” to the domain that is symbolized. Conversely, one could argue that any instance in which x ’s structural resemblance to y is exploited to find out things about y (the basic idea of simulation) is a kind of computational process that involves some form of formal symbol manipulation. In other words, formal symbol manipulation (provided that the symbols and the manipulations have been given some interpretation) really is a form of simulation, and simulation is really a kind of formal symbol manipulation.

The idea that representation involves a kind of modeling accords well with a widely shared idea about what it means for something to represent something else. Namely, the familiar idea that representations essentially “stand in” for the things that they represent. What the previous story tells us is *how* something becomes able to function as a stand-in for something else. Because of certain similarities that a representation y bears to a certain thing x , y can be used as a *substitute* for x to find out things *about* x . For in so far as certain truths about a representation follow from the configuration of properties that it also shares with the very thing that it stands for, those truths will automatically apply to that thing as well. So things can be used as stand-ins for something else, because they bear a specific kind of resemblance to those things.

Now, as I have already tried to make clear, my own thesis is not about the representational status of certain patterns of neuronal activity. Instead, my concern is with the representational status of conscious visual experience. Of course, it may turn out that conscious experience is, as a matter of fact, *identical with* those very patterns of neuronal activity. In any case, it is at least likely that certain neurons play a functional role with respect to our visual system. However, given that we are assuming that vision is direct, it seems that, in some way or other, the environment itself must actually be a constituent of the overall visual experience, in which case the neurons alone will not give us the complete story.⁶⁶

⁶⁶I am aware that there is a difficulty in the idea that something non-local to oneself is

Besides, in arguing that vision is representational, I want to appeal primarily to the *phenomenology* of vision. For it is primarily what vision is like experientially, that provides the strongest grounds for the supposition that it involves representation. Again, the representational features of vision primarily become apparent at the level of conscious experience. So it seems that questions about neuronal activity are not immediately relevant here anyway (which, I guess, should already be apparent from many of the things I have said before). In order, then, to determine whether vision is appropriately to be described as a representational process, we must decide whether vision is a kind of modeling activity. Or alternatively, whether vision is in the business of facilitating some sort of stand-in(s).

Of course, what immediately comes to mind here, is the idea that what we experience during visual perception, is really an internally generated mental model of the world. A model that is then used to reason about the world, or to find one's way in it. This would seem to be a version of the representative theory of perception. It also suggests that, whether or not vision facilitates stand-ins, depends on whether or not vision involves the generation of an internal mental model. Again however, I think that, in fact, the representational nature of vision does not primarily depend on it being indirect. Even though it may be true that in visual perception, what we directly see, is the external world itself, this does not mean that what we thus see may not in fact, in so far as it *is* seen, serve as a kind of stand-in. And as I hope to show, I think that part of the very purpose of the scenery that is made available

actually a constituent of one's experiential state. Moreover, the possibility that one could be in a hallucinatory state that is indistinguishable from "normal" experience suggests that, in fact, external properties and objects themselves never actually feature in our experience. In response to this latter problem, one could maybe offer some disjunctivist account such as can be found in (Snowdon; 1980) and (McDowell; 1986). But to be honest, I am not really a fan of disjunctivism. However, *I am* a fan of the idea that I actually see the outside world itself. And it appears that parts of that world are a certain distance away from me. And so something non-local actually appears *in* my experience, and is therefore a part of it. How that is possible, I don't know. But it does seem to be very real. It is something that I find deeply puzzling; which is why I find it so incredibly interesting.

through the process of vision is precisely to function as a kind of stand-in.

However, in order to appreciate this, it will help to first ask about the purpose of a stand-in. In other words; what is the purpose of using a stand-in rather than the thing itself? No doubt, different answers can be given here. But one important advantage of using a stand-in, I guess, is that it allows for a kind of *hypothetical* knowledge; a knowledge that is, in a sense, *independent of actual fact*. For example, oftentimes, we are interested in how some x will behave or react under certain specific conditions. The advantage of using a model, or stand-in, is that you can find out about this, without actually having to put x in the required conditions. Instead, one can use a model that, because it bears a certain similarity to x , will behave or react in those conditions exactly *like* x . And we might prefer this for various reasons. Perhaps, for example, we don't want to cause any serious damage to x (or to ourselves). Or maybe we don't want to cause any permanent change to x . More generally, perhaps we simply don't know enough about x to freely experiment with it, and so we use a model to prevent irreversible and unwanted consequences. So using a stand-in is a *less binding* (or *less definite*) way of finding out things about x (or, for that matter, the world in general).

Another reason why one might want to use a stand-in, might be that the thing itself is simply *not available*, or *absent* (which is, of course, an important component in Clark and Toribio's notion of a representation-hungry task). One could say that this is one of the primary reasons why, in mathematics, there is so much use of *symbols*. After all, due to their abstract nature, it seems that we can never actually experiment with the numbers themselves, but must perform manipulations on symbols instead.⁶⁷ But of course, things

⁶⁷This is also what makes the claim that a pocket calculator is engaged in mathematics somewhat more intelligible. An initial worry about this claim might be that it is difficult to see how a calculator device, a physical thing located at a point in space and time, can actually interact with entities that, supposedly, exist entirely outside of space and time. However, when we realize that all that the pocket calculator actually does is shuffling symbols of mathematical entities – whose symbolic status is, moreover, I claim, conferred

can also be absent simply because they are *too far away spatially*, or because, *temporally*, they *have already happened*, or *have not happened yet*. When we want to reason about things that are absent in these various ways, the use of a stand-in is a very powerful and, it would seem, indispensable tool. This is especially so when we want to reason about things and events that are currently spatially absent, or things and events that are going to happen in the future. For in these cases, it is clear that the use of a stand-in facilitates the user with *predictive* and, therefore, *anticipatory power*.

I believe this story about representation can be applied straightforwardly to vision. First of all, intuitively, there is a clear sense in which visual perceiving provides us with a kind of anticipatory power. Simply put: thanks to vision we basically are able to *avoid* crashing into things. (Of course, this is not all that vision does, but it seems to be at least one of its functions.) However, and this is important, vision does this through essentially *representational means*. In vision there is a kind of *widening* of our spatiotemporal horizon. What once was unavailable or absent to us, is now, thanks to vision, (quasi-)available or (quasi-)present. Through vision, spatially distant things are allowed to enter immediately into our conscious awareness, thereby allowing us to “reckon with” (or at least *respond to*) variables that would otherwise not have “entered into the equation”. Thanks to our vision, what was once absent, now is made present (or “quasi-present”). As we see, the extensiveness of the world itself comes into view. And with that, our own world is enlarged as well. As a consequence, we gain certain special capacities for anticipation. (Or if you want, a certain *freedom*.)

Of course, this has everything to do with the apparent action at a distance that is suggested by our visual phenomenology. But what, then, is it that makes all of this also deserving of the name ‘representational’? Well, one could say that vision basically allows us, not simply to see, but to *fore-see*.

upon them by human beings – we come to see how it is actually possible for a pocket calculator to be, *in some sense*, engaged in mathematics.

Or alternatively, one could say that viewing really is *remote viewing*. In other words, without actually having to go there, we can already know, from a distance, that something is present over there. Rather than having to go into the trouble of finding out about the thing's presence "the hard way" (by crashing into it, or by being eaten by it), vision facilitates a kind of "pre-monition" of the thing's presence "in-stead". And, of course, this is precisely why vision allows us to anticipate things. To put it a bit more metaphorically, through the process of seeing, we become cognizant of "the future that lies ahead".

Now, an objection to this idea might go as follows: the scenery that becomes available thanks to our visual apparatus cannot really be regarded as a stand-in, because the scenery is simply the world *itself*, rather than any substitute for it. (At least this is what we are supposing right now, *viz.*, that vision is direct.) In other words, suppose one is walking through the African jungle, and one sees several 100 meters away a dangerous predator. In that case, we in fact see the predator itself rather than any predator-substitute. Therefore, the predator is not really absent after all. And so there doesn't in fact seem to be any *need* for a stand-in. And thus it is dubious to suppose that visual perception is a process for the purpose of facilitating stand-ins. And therefore, one might object that vision is not really a process of representation after all. At most, vision is a process of *pre-sentation*.

I think we should disagree with this. First, note that absence is, in this context, a *relative* predicate. Although it is correct to say that the predator is not absent in the sense that it is absolutely non-existent, we are interested here in whether the predator is absent *for something else*. To be sure, what it means for something to be absent for something else may also admit of different degrees and/or interpretations. However, I think that there is a clear sense in which the things that we see at a distance are, relative to ourselves, absent. At least, *if it weren't for vision*. We naturally say that things at the other side of the earth are, relative to us, absent. But suppose

we would have the ability to see around the surface of the earth, all the way to the other side. Would it then become appropriate to say that things and events at the other side of the planet are really present to us? I don't think so. Similarly for the things we see at more ordinary distances. Those things really are, in the relevant sense, absent. It is only because we are endowed with a visual sense that their non-local presence becomes available to us.

Indeed, it is important to realize that the world that we see does not just hang in thin air. Rather, the fact that things visually appear at all, is the result of a *process*. A process, we might add, with a certain *purpose*. And, indeed, perhaps we should also add here that 'representation' is not so much a noun as it is a *verb*! The whole process of visual perception, is simultaneously a process of *making* present. A process of making distant things available for the organism, so that non-local things and events can directly enter into the factors that determine its behavior, thereby giving the organism the power of anticipation. Through the process of vision, something that is *already* present somewhere else, is now also, "over here", *made present* for the organism. And so it is indeed appropriate to speak in terms of a process of "*re*-presentation".

Moreover, this is also why we may regard the world, *in so far it is seen by us*, as functioning as a kind of stand-in or surrogate. For although we directly see the world outside of us, the fact that the world itself appears in this way *at all*, is directly dependent on a process of vision. But the *very process of vision* precisely functions as a kind of *substitute* for the fact that certain things are, in the relevant sense, absent for the organism. However, what *defines* the process of vision *as* a substitute, is basically the fact that it make (quasi-)available a world of distant things and states of affairs. A world that, in so far as it *is* made available, is indeed *inseparable from* this process. Only to the extent *that* vision makes this world available does it function as a kind of substitute process. Therefore, the world *as seen* can itself be regarded as serving a substitute role. Even though what is thus seen is the external world itself. Simply put; because there is a certain *lack*, the

process of vision basically makes available, *a world seen* in-stead.⁶⁸

Finally, we might add that, in visual experience, there is a clear, and indeed fundamental, sense in which some kind of *reference* is taking place.⁶⁹ As things appear to us in the distance, reference is being made to things that essentially exist beyond ourselves someplace else. Indeed, this power to refer seems to be intrinsic to the visual experience.⁷⁰ But of course, reference is an important characteristic of many, more uncontroversial, instances of representation; language being a primary example. As a matter of fact, what makes linguistic expressions representational is precisely the way in which they refer to other things. Linguistic expressions typically refer to things beyond themselves, often to things and places that are spatially and temporally distant. In this way, linguistic expressions make us aware of, or *make present* to our minds, things and events that would otherwise be absent to us. And this is the reason why it is appropriate to call them ‘re-presentational’. Similarly, one could say, it is precisely *because* in vision there is a reference to distant things and places that, for similar reasons, it essentially constitutes a process of representation.

So the general idea goes something like this. First, representations typically figure as a kind of substitute for something else. And they do so, because they bear a specific similarity to the things and/or states of affairs that they represent. One important feature of using a stand-in in this way, is that, in specific contexts, it allows for a kind of anticipation. Now, in vision, we are provided with a clear example of anticipation by means of representation. For what vision does, is allowing non-local things to directly “enter into the equation” of our behavior. By making us become directly aware of distant

⁶⁸This proposal is perhaps a bit similar in spirit to other theories of cognition that try to locate the representational-vehicles in the body and the world, rather than exclusively inside the brain (Clark and Chalmers; 1998), (Noë; 2006).

⁶⁹See also (Campbell; 2002).

⁷⁰Something that has, for some strange reason, escaped certain prominent skeptics of “original intentionality” (Dennett and Haugeland; 1987), (Dennett; 1987, Ch. 8).

things, vision makes present what otherwise would still be absent. This process of vision, in which non-local things are made (quasi-)present, essentially acts as a substitute for what is, relative to the organism, still missing. This is what gives vision its representational nature. (And in order to add some more intuitive force to this, we might also add that in vision, it seems essential that some sort of reference is taking place. A reference thanks to which distant places are allowed to come to our minds.)

To be sure, with respect to the points that I have made, I may only have scratched the surface of things. On the other hand, I believe that the basic idea is, really, very simple and straightforward. So perhaps there simply is, at this point, not much more to say about it. In any case, I think that the idea is clear, and indeed, quite robust. In many respects, it is clear that vision *just is* becoming aware, in a certain colorful and shapy way, of distant things. It would be difficult, to say the least, to believe that this is merely an accidental property, or side-effect, of the process of vision. Indeed, it seems more plausible to suppose that the quasi-presence of spatially distant things is among the very *purposes* of vision.

But what is perhaps more important, is that all that has been said so far about the representational status of vision neither, as far as I can see, requires that vision is indirect in the way the representative theory of perception supposes, nor requires that vision is intrinsically contentful in the way that representationism supposes. In other words, vision can be representational, even though it is direct and intrinsically contentless.

To be fair, I haven't, in this respect, said anything yet about the dispensability of content. I have only tried to show that in order for vision to be representational, it is not necessary that visual perceiving is indirect. However, I think that it is fairly easy to show that my arguments for the representational nature of vision are neither dependent on the supposition that visual experience is intrinsically contentful. This is because all the representational work is done solely by the fact that, through the process of vision, distant things get to make their appearance to us, so as to directly

enter into the behavioral equation. But the fact that things thus appear at a distance does not require that those things are also represented *as* distant. What matters is that what appears to us is, *in fact*, distant, and has, *as such*, an impact on us.

Let me try to briefly explain my point. First, if we agree that the intrinsic contentfulness of experience is indeed a substantial issue, then we must also agree that although it might be true that, on a certain occasion, one is actually looking at, say, the Taj Mahal, it is a further question whether or not the thing that one is thus visually presented with is also presented *as* the Taj Mahal. In other words, those who suppose that visual perceiving is contentless, can still coherently maintain that what appears in experience are actual chunks of the environment.⁷¹ They simply deny that being thus presented with the environment constitutes, in and of itself, *a contentful claim about* the environment. And so it is possible that, in some cases, one of those chunks is identical to, or a part of, the Taj Mahal, without it being the case that, phenomenally speaking, anything is *represented as* the Taj Mahal. Similarly, although one might actually be looking at a painting of Vermeer, it is a further question whether one's visual experience also presents the painting *as* a painting of Vermeer. And so finally, it is possible to be presented with something that is spatially distant, without it being the case that one's phenomenal experience contentfully represents it *as* distant.

But to the extent *that* we perceive distant things, a process of representation, by definition, takes place! Or at least, this is what I have tried to argue. For to make things that are present somewhere else also (quasi-)present "over here" *just is* to re-present. In other words, vision is representational, because it facilitates an awareness of spatially distant things, not because it is in the business of constructing a representation that proclaims *that* such and such

⁷¹Of course, it is also possible to suppose that perception is contentless and *deny* that the phenomenal aspects that are given in it are identical to parts of the external environment. In this way, one might end up with various views. For example, both some version of vehicle-based representationalism or an adverbial theory of phenomenal experience ((Fish; 2010, Ch. 3)) seem to be possibilities in this respect.

is over there in the distance. And one might add that although, plausibly, the function of this process of representation is to make distant things become, in a certain way, causally relevant to the behavior of the organism; for this, also no content is required. Indeed, this causal relevance seems to be realized as soon as distant things actually start to figure in our sensory experience. For as soon as they do so, they, by definition, are making an impact on us. And plausibly, we are wired in such a way that, in those circumstance, we are automatically moved to respond in appropriate ways. Indeed, it may very well be that our responses to distant things are entirely “dumb”, “instinctive”, and “unreflective”. What matters is that *what* we respond to are *distant* things, and that the sense of vision facilitates that such responding is even possible.

In any case, I hope that my point is clear: the kind of representational activity that vision is engaged in, is not only not dependent on whether or not vision is indirect, it is neither dependent on whether visual perceiving is contentful. We only need to consider vision in its bare naked form, to find that these issues have little to do with the representational activity that seems so characteristic of vision.

IV

I believe that I have succeeded in showing that there is a straightforward sense in which vision is a process of representation. And so I have reached the main goal of this paper. However, one might object that, although there may be *some sense* in which it is appropriate to refer to vision as a process of ‘representation’, the sort of representation that one ends up with is not of a very interesting kind. Not very interesting, precisely because it does not involve the claim that we only perceive the world indirectly, nor the claim that perceptual experience is contentful. And it is indirect perception, and even more so, mental content, that philosophers have been worrying about.

Recent ‘non-representationalist’ theories, most notably radical enactivism,

are primarily concerned with showing the dispensability of *content*. And so it would seem that my arguments for the thesis that vision is representational, do not actually show that non-representational theories are, in any substantial way, inadequate. After all, it would seem that these ‘non-representationalist’ frameworks are reacting against a notion of representation that is essentially different from mine. Thus, someone might point out: “If you twist long enough the meanings of our concept of representation and vision then, surely, you will find some way to connect the two. However, it seems that all you have accomplished this way is a mere terminological victory, not something that is of any substantial relevance to the current debate on representation. After all, you have not actually established the fundamental importance of ‘mental representation’ in the truly relevant sense of that term. Your arguments are beside the point!”

I beg to differ. If vision is indeed as I have described it, then it seems that we still have to face certain important problems. Problems that are essentially the same as the those that come up with the more traditional conceptions of mental representation. For example, although it might seem that eliminating mental content saves us from the problem of how any such items as mental contents can have a true causal impact on the world, a problem that is essentially similar remains.

As I have pointed out in the opening section of this paper, there appears to be, on the classical cognitivist account, a certain tension between the idea that the brain is merely a syntax-sensitive processor, and the phenomenological observation that our mental states appear to be transparent to us in such a way that we seem to be able to directly intuit the very *contents* of our thoughts. The tension here arises from the apparent fact that, contrary to what the materialist-computational view implies, mental content *does* seem to be directly involved in shaping how we think and behave. At the very least, if it is true that we are merely syntax-sensitive machines, then it sure seems strange that this illusion of transparency is nevertheless there.

It would seem that disposing of content as a mere illusion would solve

the tension here. If phenomenal states are not intrinsically contentful, then we need not worry about the causal efficacy of content. However, even if we drain phenomenal states from all content, it seems right to say that those states nevertheless retain their transparent nature.⁷² Especially so in the case of vision. For, even though its phenomenal deliverances do not involve any contentful claiming, this does not make it any less true that in vision we primarily look out into the world. We do not in a similar way “see” the internal state that *brings about* this looking out. When we try to introspect the visual experience itself, it seems that there is nothing there that is phenomenally manifest, other than the world itself that is already beyond our skins. Disposing of content doesn’t really change much about this. But the kind of transparency that one thus ends up with is equally puzzling and problematic.

For the puzzling part, consider what the philosopher Howard Robinson has to say about what he refers to as “...one of the basic problems of perception”⁷³:

On the one hand, a little reflection – that is, thought that does not resort to any science that goes beyond common experience – shows that perception involves some sort of physical influence running from the external object to the sense organ of the perceiver. On the other hand, the essential nature of experience seems to be that the subject mentally reaches out to, and makes conscious contact with, the external object. The directionality of the physical process and that of the lived experience seem to be in direct conflict. How can a process in which the subject is the passive recipient of a stimulus be the physical aspect or realisation of a process in which the subject reaches actively and consciously out into the world? (Robinson; 1994, p. 4)

I don’t think that I could have formulated it better myself. There is some-

⁷²See also (Kennedy; 2009) for a somewhat related conclusion.

⁷³See (Robinson; 1994, p. 4).

thing counter-intuitive about the physical “inwardness” of the physical stimulus and the outward direction of our mental regard as we see something in the distance. And this is mainly because we suppose that the two are in fact two different aspects of one and the same process.⁷⁴ How is it that the kind of visual processing that happens inside of us when incoming electrochemical signals excite our retina’s, contains *within itself* a reference *outwards* to something that is located far beyond this processing?

To be sure, what we have here is not mere counter-intuition and puzzlement. I believe that if we truly think about what appears to be going on here, we come to see that there are some very real problems to be dealt with. Of course, what we are immediately reminded of here, is the hard problem of consciousness. How do we account for the fact that we consciously experience ourselves and the world, in terms of (physical) processes that are not themselves conscious? How do you get something conscious from something non-conscious? As I have already confessed, I don’t think that we have the answer to this yet. However, the problem here is not simply why it is that our internal states are “charged” with a “phenomenal glow”; the problem is *deeper*. The phenomenology of vision also seems to violate our basic causal and geometrical intuitions about the physical world.

It is not the mere *phenomenal* side (in the technical sense of that term) of vision that is problematic, but also the fact that, on the face of it, vision appears to allow non-local factors to directly enter into the set of variables that shape our thinking and doing. The difficulty here is not just how to accommodate a place for the phenomenal in the physical; it is also very much a problem of how to reconcile problematic forms of interaction with our current physicalist metaphysics.

On the face of it, vision seems to involve *instantaneous action at a distance*. As we see things at various distances, an immediate interaction be-

⁷⁴We are a bit reminded here of the so-called “wagon-wheel effect”, a perceptual illusion in which a wheel appears to rotate contrary to the direction in which it actually rotates. Of course, whether visual phenomenology is equally illusory remains to be seen.

tween two distant things appears to take place. A kind of interaction that essentially appears to be non-local. However, given our general understanding of the physical world, we feel that such kinds of interaction ought not to be possible. At least, I think it is fairly uncontroversial that most of our physicalist theorizing proceeds on the assumption that all interaction must ultimately resolve itself in contiguous micro-physical causal interactions.⁷⁵ Therefore, visual phenomenology appears to contradict what we suppose the physical world is like. And of course, from a “physicalistic” or “naturalistic” point of view, this is bad news (unless we are willing to suppose that the natural contains in fact more than we thought it did).

Given that we are here, apparently, confronted with a contradiction in our conceptions about what the world is like, there are several ways to respond. And given the metaphysical climate of our times, it seems that the most obvious response is try to show that the phenomenological data are in fact *merely apparent*. However, it seems that there are then two further options. One can either go for a *reductionistic* approach, and try to show that our awareness of the distant object is in fact wholly reducible, without remainder, to local physical interactions at a lower ontological level, or one can adopt an *eliminativist* strategy, and try to show that the phenomenological data are illusory, such that we are not actually aware of any distant entities or states of affairs. So the reductionist strategy is to *accept* that we *are* aware of

⁷⁵I am aware of the debates among theoretical physicists, and philosophers of physics, about the question of whether the ‘principle of locality’ – according to which physical things and processes can only be influenced by their immediate surroundings – is actually true of our world (Berkovitz; 2014). But I am going to skip that difficulty for now, since, much to my regret, I don’t know enough about physics. I think it is safe to say that the principle of locality is an implicit assumption in much of the physicalist literature (The Philosopher C.D. Broad, for example, has referred to it as one of the “Basic Limiting Principles” (Broad; 1949). See also (Griffin; 1993).) At the very least, I think it is fair to say that, when it comes to the direct causal determinants of human behavior, the default position among naturalistically-minded theorists, is that only things that are internal to us, or physical forces that directly impinge on our bodily surfaces, are direct causes of our behavior.

something in the distance, but to *deny* that it involves any actual unmediated action at a distance. The eliminativist strategy, on the other hand, is to *deny* the phenomenological data *altogether*, and thereby simply to forestall the need to account for the alleged anomalies. Unfortunately, I think that both options are problematic.

In what follows I would like to go into this a bit. However, what I am going to say will be rather schematic, since, at this point, it is not possible to delve into new problems in any sufficient detail. The main reason that I want to touch upon these issues is that, as we have seen, some of the main problems with mental representation precisely have to do with how they are to have a causal impact on the world. And I think it must be pointed out that, even if we deny that (basic) mentality involves the more traditional kinds of mental representation, we are still left with a serious interaction problem. An interaction problem that basically comes about because visual perception apparently allows us to directly engage with non-local entities (and which is, as I have argued, precisely what makes vision deserving of the name ‘representational’). Of course, it remains to be seen if no reductionistic or eliminativistic account can actually be given here. But if it can’t (and I do think that both reductionism and eliminativism are problematic) then it does seem to me that many of the new non-representationalist accounts largely leave the main problems untouched.

Let me start with eliminativism. If we choose to be eliminativists about the phenomenology of vision, then again, there seem to be various options. First of all, we might say that, rather than being immediately aware of something distant and non-local, what we are actually immediately aware of is *something else*, something *local* and *internal* to us. This leaves one with some version of the representative theory of perception; vehicle-based representationalism. Not, I suppose, a very desirable outcome.⁷⁶ Another possibility is to say that,

⁷⁶Why? Well, first of all, it just doesn’t seem to be “true to the phenomenology”. As I have pointed out, experiences seem to have a kind of transparency to them. What is phe-

rather than being aware of an external object, one is actually aware of an intentional content. In effect, one ends up with some version of content-based representationism. Choosing the latter option eliminates the difficulty of accounting for how a spatially distant object can directly affect us in the way visual phenomenology suggests. (For according to a widely held view, intentional contents are not “object-involving.”) However, this leaves one, again, with the problem of how it is possible for *intentional contents* to have an impact on the world. In any case, it seems that one ends up with either one of the more traditional representational theories of visual perception. And it is doubtful that this actually solves any of our problems.

Are there any other eliminativist options? Are there other ways to deny that in seeing we are actually confronted with distant chunks of the environment? It is difficult to think of any concrete examples, but I guess the closest thing that comes to mind is a kind of dispositional account à la Daniel Dennett.⁷⁷

Unfortunately, I cannot give a full treatment here of Dennett’s views. In fact, I am going to be rather brief in my criticism with respect to the sort of proposal that Dennett has on offer (or at least my understanding of it).⁷⁸ Admittedly, I am also not sure if Dennett has ever been explicitly concerned with our current problem. Especially my own formulation of it. However, I

nomenally manifest to us is what the experience is *of*, not the underlying vehicular process that supports it. But vehicle-based representationalism is also undesirable for ontological (due to its apparent commitment to the phenomenal principle) and epistemological reasons (it sets up the dreaded veil of perception).

⁷⁷See, e.g., (Dennett; 1988), (Dennett; 1991).

⁷⁸I have considered including the account of George Rey here (Rey; 1983), but I am not sure if he really belongs. This is because, on the one hand, he seems to be an eliminativist about consciousness, whereas on the other hand, he appears to be a realist about representational content. However, he also seems to think that representational states are straightforwardly reducible to the kind of materialistic computational processes of classical cognitivism. And, therefore, it is doubtful if his representational states actually do the required representational work. Anyway, I think Rey’s account is close in spirit to Dennett’s. And my criticism of Dennett’s view therefore also possibly applies to Rey’s.

think it is safe to say that, according to his line of thinking, the response to our problem ought to go something like the following.

From a Dennettian point of view, we might say that, although it *seems* as though we are actually aware of distant things in the outside world, we are not *really* so aware; instead, we only *think* that we are directly aware of the outside world. The latter of which simply reduces to being disposed to behave and talk in a certain way.⁷⁹ In other words, the world itself doesn't *really* appear to you, you only *think* it does!

Let me briefly say something about why I think this view is problematic. If this line of thinking entails that we are not actually consciously open to the (distant) world, then my main problem with this view, is that it seems to lead to a certain kind of skepticism, or a-gnosticism. For if what I see before me is not actually the world itself, but is rather, somehow, a symptom of my own dispositions, *then I simply don't know what "world" is anymore*. The consequence of this dispositional view seems to be that the term 'world' itself loses its very reference-base. When it comes to the reference-fixation of the term 'world', it seems that some causal theory of reference is essentially correct. After all, the fact that one even *has* a notion of the fact that there is any such thing as "the world" *at all* is, first and foremost, grounded in conscious experience (recall the thing's I have said at the beginning of part II). However, if what is phenomenally manifest to me in experience, and what I have always *ostensibly referred to* as 'the world', is *not* in fact the world, then I guess I simply don't know anymore what "world" signifies. (If *that* is

⁷⁹That this would be an appropriate Dennettian response is perhaps best indicated by the following fragment from a section of *Consciousness Explained* called 'Seeing is Believing': "Now you've done it. You've fallen in a trap, along with a lot of others. You seem to think that there's a difference between thinking (judging, deciding, being of heartfelt opinion that) something seems pink to you and something *really seeming* pink to you. But there is no difference. There is no such phenomenon as really seeming – over and above the phenomenon of judging in one way or another that something is the case. ... There *seems to be phenomenology*...But it does not follow from this undeniable, universally attested fact that *there really is phenomenology*" (Dennett; 1991, p. 364 – 366) (emphasis from the original text).

not the world, then I don't know what is!) In effect, my world and some of my most basic concepts are shattered, and it is indeed doubtful whether any further discourse is even possible.

Or equally worse, on Dennett's account, it would seem that the world is merely what I *think* or *judge* it is! At the very least, Dennett seems to suggest that, phenomenologically speaking, there is nothing aside from my own judgments. But it is difficult to see how this can be anything other than cartesianism (or even idealism). An ironic consequence, given Dennett's philosophical motivations! In any case, it renders deeply problematic how I, and every one else, can have any basis for knowing what the world is actually like if, phenomenologically speaking, everything is determined by one's judgments. Yet I'd like to think that what appears to me in experience, is at least partly *anchored* in something that exists *outside* and *beyond* me. If that is not the case, then the very possibility of objectivity appears to fall apart.⁸⁰ A price, I believe, that is too high to pay.⁸¹

Admittedly, I don't know if this criticism is entirely fair to Dennett. Maybe I am oversimplifying his views, or maybe I am merely setting up a straw man here. But then again, it is often very difficult to discern just what exactly Dennett's views *are*. In addition, I guess it is also just very difficult to say what being an eliminativist about our apparent openness to the (distant) world amounts to! Really, what does it mean to say that, most of the time, you are not really seeing distant events and objects? Of course, the possibility of hallucination and dreaming incredibly complicates everything. Indeed, if such things as hallucinations and dreams simply did not exist, many of the most intricate problems of philosophy would simply not arise! And I am

⁸⁰For what I consider to be a similar point, see (Kim; 2000, p. 31).

⁸¹Not to mention the fact that such universal subjectivism is ultimately self-undermining. Note also that appealing to some collective process of triangulation does not seem to help much either, for all that one ends up with is simply more judgment. And who knows that we are not perpetually in a collective state of mass-delusion?. Or worse, if all that we have access to are our own judgments, then what grounds our judgment that there is even a world at all?

aware that, in the present context, it would be good to say a little more about the implications of hallucinations and dreams, especially given my present commitment to the idea that perception is direct. I am not going to that, however. I guess that, ultimately, the supposition that we are genuinely confronted with something outside of us, is something that is non-negotiable. Whatever else is true, there is something “out there”, outside and beyond me. And I am looking at it! So although my arguments may be wanting, I am simply going to assume that eliminativism about the phenomenology of vision is wrong (or at least, not clearly defined). Let me, therefore, turn to reductionism, which is where the really interesting problems are.

A reductionist account of the action at a distance that appears to characterize visual phenomenology grants that the phenomenon is pretty much like we think it is, but denies that it involves any problematic sort of interaction. In other words, it denies that the apparent action at a distance, that seems to take place when we see, is in any way *instantaneous* (or ‘immediate’, or ‘fundamental’). Instead, it supposes that the phenomenon is ultimately reducible, without remainder, to micro-physical interactions of an intuitively unproblematic sort. Although it is true that we are visually aware of distant things and events, this awareness, and the kind of interaction that it seems to afford, is entirely reducible to causal chains of contiguous physical interactions that take place at the microscopic base-level of existence.

On the face of it, reductionism seems to be both plausible and desirable. After all, it respects the phenomenological data, and it is in accordance with how we suppose the physical world is like. Furthermore, it seems just obvious that visual perception cannot be instantaneous action at distance, because we know that visual perception is thoroughly mediated by physical and physiological processes. It is common knowledge among most educated people that seeing a distant object depends on a light source. Our seeing of a distant object is mediated by the electromagnetic waves that are reflected by the object, and that *reach out to us* and stimulate our retina’s. (Of course,

a lot more is going on, but the point should be clear.) Without this whole process, there would simply be no visual perceiving. So it seems naïve and outright false to suppose that there literally is a direct influence between us and the distant object. Visual perception does not involve instantaneous action at a distance, and therefore it does not threaten our basic naturalistic assumptions about how the world works.

To be sure, so far, no one seems to have come up with an actual account of how we get from the physiological and physical details, to the event in which the phenomenal grandeur and vastness of the world is allowed to come to the fore, as part of a visual experience. And indeed, it is very difficult to conceive of what such an account might look like. But this does not mean that visual perceiving does not in fact exhaustively reduce to micro-physical processes in the relevant way. And given the metaphysical stakes, it seems therefore best to work with the assumption of reductionism.

Although I'm afraid that I won't be able to give any knock-down arguments, I do want to formulate some concerns that I have about this reductionist proposal. For, intuitively, I feel that things don't "add up" in quite the way that reductionism hopes. In any case, if reductionism *is* true, then it just doesn't make sense *why we see at all*. If reductionism is true, then there simply does not appear to be any *purpose* to the fact that we are actually visually aware of things in the distance. And since I believe this is a rather unlikely proposition, perhaps it is better to simply give up reductionism, and think about new ways in which the different layers of existence might be related. Let me explain.

First of all, let us reflect on the reductionist supposition that, in reality, the only things that exercise a direct causal influence on us are physical forces that are local to us in space and time. In the case of vision, the relevant physical forces seem to be, mainly, the electromagnetic waves that stimulate our retina's, and those that are implied by the further, internal processes, that are thereby engaged. But consider now the following point: if the electromagnetic stimuli and internal processing do all the causal work, then

why should *we* also, *in addition*, see the distant object as well? Whatever the answer, there is a strong temptation to say that, really, this is a very bad case of causal over-determination. In any case, if the only things that can have a direct causal impact must be physical processes that are immediately local to us then, intuitively, it just doesn't make sense why we actually perceive non-local things and events in the distance. Of course, the problem here is essentially similar to the problem of the causal potency of content. Consider the following comment by Jaegwon Kim:

Thus, that a given intentional state of an organism instantiates a certain semantic property is a *relational* fact, a fact that essentially involves the organism's relationship to various external environmental and historical factors...If inner states are implicated in behavior causation, it seems that all the causal work is done by their "syntactic", or at any rate internal/intrinsic properties, leaving their semantic properties causally otiose. The problem of mental causation generated by syntacticalism therefore is to answer the following question: *How can extrinsic, relational properties be causally efficacious in behavior production?*

So the crux of the problem lies in the supposed fact that mental properties, in particular, content properties...are relational properties, extrinsic to the organisms instantiating them, whereas we expect the causative properties of behavior to be intrinsic and internal. (Kim; 2000, p. 37)(Emphasis from the original text)⁸²

It should be clear that something very similar is going on in visual perception. If only local causal forces are relevant to behavior, it would seem that the distant object that we have in view in conscious experience, does not actually have any direct causal influence on us in bringing about thought and behavior. But then what is the point of seeing things in the distance? Why

⁸²See also (Fodor; 1994) for an extensive treatment of this issue.

does one see at all? Of course, in the previous case, we can simply *eliminate* the notion of content from our theorizing. However, we cannot as easily eliminate from our theorizing the fact that we seem to see non-local things and events at various distances. Even if we drain phenomenal experience from all intrinsic contentfulness, the distant environment remains in view nonetheless. Contrary to what the radical enactivists claim, supposing that phenomenal experience is contentless does not actually solve the problems.

Some might object to all this that, in the previous kind of reasoning, the problems arise because there is too much of an exclusive focus on sensory stimulation. But everyone agrees that mere sensory stimulation is not enough for full-blown perception. For this, something *more* is needed. There are various theories about what this extra ingredient might be. There seem to be roughly two kinds of proposals.⁸³ According to the more traditional view, what is needed are extra representational resources that are native to the brain itself.⁸⁴ More recent theories however, such as radical enactivism, typically hold that, in order to get full-blown experience, sensation must be supplemented by temporally extended dynamical interactions between the organism and the environment.⁸⁵ In any case, mere sensory stimulation alone is not enough. Something extra is needed for actual perceiving.

However, I think that this criticism is besides the point. Of course, it is true that a lot more processes beside retinal stimulation are involved in producing, or realizing, visual experience. But as long as these processes are internal to ourselves, or exclusively take place at the interface of our bodily surfaces, nothing substantial is gained. And as far as I can discern, the kind of processes that the aforementioned theories appeal to, all seem to fall in this category. In any case, the picture that we are left with is still that of an organism whose behavior is only directly causally influenced by local or

⁸³See (Hutto; 2011).

⁸⁴A kind of Kantian strategy. See e.g. (Fodor; 1975, Ch. 1).

⁸⁵See (Hutto and Myin; 2013, P. 158); a way of thinking that seems much inspired by Gibson's ecological theory of perception (Gibson; 1972). See also (Chemero; 2009) for a similar approach.

internal factors. Moreover, appealing to representational contents that are native to the brain would, in the present context, not be an option anyway. For these either boil down to mere syntactic force, which makes their status as representations immediately questionable, or they involve causally efficacious contents, which is, as should be clear by now, equally problematic, especially if we want to be reductionists.

So the only option left seems to be something like radical enactivism, according to which intelligent, perception-guided, behavior is a matter of being “information-sensitive” in the right way. However, I sincerely doubt that the radical enactivists have the right tools to deal with phenomenal experiences, especially those of the visual type. This is primarily because they seem to want to account for phenomenal experience entirely in terms of the idea that basic mentality, including our basic modes of perception, is ultimately a matter of “information-sensitive responding”.⁸⁶ Although, from an empirical-physiological point of view there is certainly much to be said for this, it seems to me that the language of “responding” is inadequate to capture or ground phenomenal experience. Indeed, I would say that things are rather the other way around: many of our intelligent respondings are precisely grounded in, or enabled by, (visual) *phenomenal experience*. Intuitively, phenomenal experience, characterized as it is by world-openness, *enlarges* the spatiotemporal confines of the environment, and seems thereby to facilitate a more extensive range of environmental offerings immediately there for us to respond to. So rather than being itself wholly *constituted by* “respondings”, it would seem that phenomenal experience rather *allows* for certain new types of responding to *emerge*.

To be sure, Hutto and Myin do speak of mentality as an “emergent” phenomenon. For example, somewhere they state “Sentience and sapience *emerge* through repeated processes of organismic engagement with environmental offerings”.⁸⁷ However, I doubt that Hutto and Myin use the term

⁸⁶See (Hutto and Myin; 2013, Ch. 4).

⁸⁷See (Hutto and Myin; 2013, p. 8) (emphasis added by myself).

‘emergence’ in its proper technical sense; the sense in which it denotes an irreducible high-level phenomenon with new causal powers of its own (causal powers that are irreducible to the sum of causal powers that are present in its emergence-base). Given their naturalistic motivations, it is unlikely that they actually endorse the thesis that mentality is a genuinely emergent phenomenon. However, this does bring me to an important point; intuitively, vision is supposed to allow the organism with new ways to interact with the environment, but it seems that reductionism is unable to give an account of how this could be so.

But let me turn to another consideration first. If the reductionist is right, then the only immediate causal determinants of our thoughts and behavior are localized internally within us, or at the immediate surface of our bodies. But if this is true, what to say, then, about so-called “phenomenal judgments”? For they would seem to be, in some important sense, illusory.⁸⁸

Look around you, and direct your attention to a reasonably distant object in your visual field. Notice how it seems, in its own way, present to you. Pay close attention to the robustness of the fact that something in the distance seems to fill up your visual field. As you do this, notice that you also have the seeming ability to consciously reflect on the distant object in all its “suchness”; that you can in various way let your thoughts go over the object (“*what is that thing?*”) Indeed, notice that the object itself seems to be the immediate object of your thinking. There seems to be no mistake about it: *because* the object is there, showing itself in plain view, you can consciously reflect on it. In other words, there is a strong sense in which there is a special – admittedly difficult to describe – intimate immediacy between the object that one sees and the judgments that one forms about it.

Yet if, as reductionism holds, the direct causes of your own thinking processes are in fact entirely local and internal to you, then it would seem that, in some strange way, all this is merely a kind of illusion. At the very

⁸⁸See also (Chalmers; 1996, Ch. 5).

least, it seems that the causal relation between the seeing that happens “over here”, and the object seen that is “over there”, is not as it evidently appears to be. As a matter of fact, not only would it seem that you are in a fundamental way mistaken about the true cause of your *judgments*, it would even seem that you are mistaken about the true cause of your very *awareness of the distant object itself*. For in so far as your apparent *awareness of the distant object* is rather an immediate result of what happens to you “over here”; what phenomenally appears “over there” cannot actually be any direct cause of what happens “over here”. And so perhaps we may even correctly conclude from this that you are not aware of anything distant at all.⁸⁹ But, reductionism precisely wants to save the phenomena! So this result seems contrary to the reductionist motives. Does this mean that you cannot both respect the phenomenological data *and* be a reductionist at the same time?

The reductionist might respond that the previous problems arise because we mistakenly suppose that awareness-at-a-distance happens at the micro-physical level. In other words, the mistake is that we are looking for awareness-at-a-distance at an ontological level that simply is too low. One could argue that this is similar to trying to find representation at any of the sub-parts of an organism. It is only at the level of the organism *as a whole*, that it is appropriate to speak of something being visually cognizant of distant things and events; it cannot be found at lower levels of description. In other words, the reductionist might suppose that, although no individual component of the reduction-base is in the relevant way cognizant of the distant environment, the reduction-base, *taken as a whole*, is. After all, it seems perfectly coherent to suppose that different micro-physical causes together may form a macro-cause, and that the various micro-physical effects that they bring about together form a macro-effect. Similarly, reductionists may suppose that awareness-at-a-distance is a causal phenomenon that takes place

⁸⁹Not, at least, in so far as being aware of something is a causal relation, of some sort, between the act and object of awareness. But let us not go into any further complications here.

at a macroscopic level. So although, at the micro-physical level, nothing is directly influenced by the distant object that figures in visual perception, the organism as a whole is.

This proposal is also, I believe, problematic. First of all, although we might suppose that it is in fact the organism as a whole, rather than any of its sub-parts, that is visually cognizant of the distant object, there still remains a *spatial gap* between the micro-physical processes that together constitute the organism as a whole, on the one hand, and the distant object that is seen, on the other. So we are still confronted with the problem of how a *macro-physical* entity can be visually influenced, in the relevant way, by another physical entity that is distant to it. It seems that, in order for the distant object to make any sort of causal relevance to the macro-physical entity, some sort of physical force must actually cross the gap and reach our bodily surfaces. But then, again, we seem to end up with the problem of *why* the distant object should enter our awareness *at all*, since all the relevant causal work is still entirely local to us.

Of course, we may speculate that, perhaps the *wave-particle duality* of light can give us the further resources that are needed to account for the phenomenon of awareness-at-a-distance in vision. After all, if light is what mediates the visual connection, by crossing the gap between the seer and the seen, and light is also both field-like *and* particle-like, then perhaps, if we exercise our imaginations long enough, we may come to see how something that reaches out to us can nevertheless contain some sort of reference to its source. Although this is an interesting idea, and although there may be yet other ways to “widen” the reduction-base, so as to cross the spatial gap, it seems to me that reductionism, in any case, faces a more general difficulty.

From the reductionist perspective, our being visually aware of distant things, ultimately reduces, without remainder, to a set of low-level micro-physical interactions. On this view, our visual-awareness-at-a-distance is a macro-causal phenomenon, whose causal powers are *nothing over and above* the sum

of causal powers that belong to its reduction-base components. Now although there does not appear to be any strict logical impossibility in this view; it does lead to very counter-intuitive results. The problems here precisely have to do with the reductionist view that the causal powers of macro-phenomena are nothing over and above the sum of causal powers of their reduction-base components. This is counter-intuitive because, in our case, it would seem that vision does not have any causal properties of its own. And so it would seem that the macro-phenomenon of vision, *qua* vision, does not have any true causal impact on the world. Again, the result is that it renders unintelligible why we actually see the (distant) world at all.

To be sure, my concern here is not that, from a reductionist perspective, there really is no such thing as genuine macro-causation.⁹⁰ Very clearly, there *is* macro-level causation. Even for the reductionist. It's just that macro-level causation is exhaustively constituted by instances of micro-level causation. Just because a table is nothing over and above the sub-atomic particles out of which it is made, does not mean that there is no table. A similar thing holds for macro-level causation. The problem, rather, lies in the fact that, apparently, some macro-physical processes are in fact *subjects of phenomenal states*, whereas others are not. And, from a reductionist perspective, there really does not seem to be any *rationale* for why this should be so. Let me try to explain why.

First of all, it seems correct to suppose that, when it comes to giving a description of reality in macro-level terms, there is actually more than one way to do this, and that, necessarily, there is a kind of arbitrariness involved when it comes to which one of these ways we choose. This arbitrariness seems to result from the fact that describing the world in macro-level terms, requires that we "*carve-up*" the world in a certain way, whereas nature's own "*natural joints*" only seem to be objectively pronounced at the microscopic level. Of course, we might say that there do seem to be orderly and less orderly ways of dividing up the world into macro-level phenomena. And

⁹⁰See also (Kim; 2000, 77-80).

indeed, some ways of articulating a macro-physical world seem to come more natural than others. However, it would seem that the kind of patterns that are, in this way, more salient to us than others, is more a function of our own cognitive and perceptual apparatus than an actual difference that is out there in the world.⁹¹ Of course, this is not to say that the kind of patterns that are salient to us do not actually exist. It's just that their salience is in a sense agent-relative. From the cognitive point of view of a creature that is wholly different from us, other patterns and orders might naturally be more pronounced.

In any case, let's assume that there is a great variety of ways to carve up and articulate a macroscopic world, and that any one of them is, in a way, as good or arbitrary as the other. Although we might divide the world into different macroscopic "units"; objectively, these ways of drawing the boundaries are merely optional.

If we proceed further on the basis of this supposition, we must also say that each way of articulating or carving out a macro-causal event is as good as any other other. After all, each macro-causal event is nothing over and above the sum of micro-causes and micro-effects to which it reduces. And so, in this respect, they are really all the same. At the macroscopic level, we shouldn't expect any fundamental (meta)physical differences between various macro-causal events. In the end, they all boil down to various sets of microscopic interactions. Of course, we may favor some causal divisions over others, but this seems to be primarily for practical reasons, rather than any metaphysical reasons. Ultimately, all macro-level events seem to be on a par.

The thing with this whole story is that, although it sounds reasonable, it does not seem to be entirely true. For contrary to our previous supposition, it *does* seem to be the case that at least *some* macroscopic "carvings" are naturally, in and of themselves, more pronounced than others. For notice that, as human beings, *we ourselves* seem to have our existence at a very

⁹¹Consider, for example, the difference between the way a language sounds to its native speaker, and the way it sounds to a foreigner.

specific macro-level of reality. After all, our bodies are macro-level entities. But what seems to set them apart from other macro-level phenomena, is that they have a very pronounced, non-arbitrary, natural unity to them. A unity that is essentially defined by the fact that each human living body is a *subject of experience*. Intuitively, not every arbitrarily carved-out macro-physical process has such a special subjective unity to it. Rather, it seems that only certain rather specific macro-level defined units are also subjects of phenomenal experience. The problem, however, is that from a reductionist point of view, it seems difficult to account for why exactly phenomenal subjects are distributed in the particular way that they are at the macro-level, and not in some other way.

Again, recall that any way of carving out a macro-causal event seems as good as any other. They are just ever so many different ways of defining a set of micro-causal events. Objectively speaking, no macro-causal events seem to be more “unified” than others. It’s all just a matter of how we carve up the infinite ocean of microscopic events that is our universe, into larger macroscopic structures. Yet, it seems that, in some special cases, some macro-level carving-outs do seem to be substantially different from others. For it seems that the set of microscopic processes that these macro-carvings pick out appear to be characterized by special sort of unity; a subjective unity that is expressed by the fact that the microscopic processes that they pick out together constitute a subject of experience.

However, the difficulty with a reductionist approach, is that these macro-subjective unities don’t have any special causal status. Everything that these macro-level structures bring about can be entirely accounted for in terms of micro-physical causes and effects. In other words, although certain micro-physical processes are subjectively unified in a seemingly substantial way; it appears that this clear-cut macro-level individuality simply does not *do* anything. It doesn’t leave any extra traces on the world, except, apparently, logically private ones. But why should macro-level individuals come into existence if, as it seems, they don’t do anything extra? If everything is

exhaustively accomplished by micro-physical processes, then why should any macro-level subjects appear to keep a check on things? Why should there be macro-phenomenal worlds manifest at all? There just not seem to be any rationale for it.

Of course, these “vertical” considerations do not in and of themselves constitute a knock-down argument against reductionism. Perhaps it is just a brute fact that some macro-level phenomena are subjects of experience whereas others are not, everything else being the same. However, if we also add here the “horizontal” problems that I considered earlier, it becomes increasingly clear that there is something fishy about reductionism. Especially in the case of vision. If we consider how microscopic and the macroscopic relate to each other “vertically”, there does not seem to be any rationale for why, somewhere along the way, aggregates of micro-physical processes should be subjects of visual experience. And if we consider how the seer and the seen are causally related “horizontally”, we are again faced with the problem of why there is seeing at all, given that the only real causal contributions of vision must happen locally (i.e., internally and at the sensory surface). Of course, this still does not constitute any effective knock-down of reductionism. But still, it invites difficult questions that need to be answered.

What are the alternatives to reductionism? Intuitively, we may suppose that visual-awareness-at-a-distance is really an *emergent* phenomenon. One that brings into existence new causal powers of its own. The intuitive force of emergentism with respect to vision stems from the plausible idea that vision primarily *makes possible* our ability to directly respond to the distant environment, rather than that it already itself *constitutes* such responding. It seems natural to suppose that only *after* vision has brought the distant environment into view, can we effectively respond to and anticipate it. The process of vision allows something that is external to that process to come into view (this is what gives it its transparent character). However, it would be counter-intuitive to suppose that this is an end in itself, and that the

causal chain stops there. It seems more natural to suppose that the visual process brings the distant object into view *in order that* the organism can directly anticipate it from a distance, providing the organism with a kind of premonition power. In other words, the process of vision seems to facilitate a *further* and *direct* causal connection between the organism and the distant object, that is *over and above*, and irreducible to, the micro-physical interactions from which it emerges.

Unfortunately, it is unclear whether such emergence is metaphysically coherent. For it implies that there is some kind of downward causation from macro-level phenomena to micro-level phenomena. And, as some have pointed out, the idea of downward causation is problematic.⁹² A better alternative, perhaps, is *holism*, which I would define as the view that, somehow, all the various ontological strata, from the microscopic all the way up to the macroscopic, are already necessarily co-occurrent, in such a way that neither one of them exclusively constrains all the others. In other words, it holds that it is wrong to suppose that the microscopic has ontological precedence over the macroscopic, and that it is better to say that microscopic phenomena are always already subject to macroscopic constraints. This allows one to say, for example, that smaller spatiotemporal slices are perhaps already constrained by larger spatiotemporal slices (something that may perhaps already be implied by certain “eternalist” conceptions of time). Of course, all this is not to say that everything is entirely constrained from the top down. Perhaps the implication here is that ‘bottom-up’ and ‘top-down’ are, ultimately, misleading notions when it comes to the true nature of reality.

Finally, it is perhaps worthwhile to point out, again, that the principle of locality is not an uncontroversial premise anymore among theoretical physicists and philosophers of physics. Certain results in quantum physics may suggest that instantaneous action at a distance may in fact already take place at the micro-physical level of reality.⁹³ And again, the wave-particle duality

⁹²See (Kim; 2006).

⁹³See (Berkovitz; 2014).

of light may suggest new ways to understand the phenomenology of vision. For note that theoretical physicists nowadays usually seem to refrain from giving any actual physical interpretations of their theories. But perhaps it is interesting to play with the idea that the phenomenology of vision, and the intrinsic properties of light, can maybe shed light on one another. Moreover, as some have noted, the kind of mechanistic view of physics on which most of our current philosophical theorizing seems to be built, is really in many respects “oversimplified” and out-dated. As David Bilodeau notes “As Physics has developed, it has been necessary to include concepts which are non-geometric in the Cartesian sense. A well-known example is the difficulty of formulating a mechanical description of the electromagnetic field (i.e. describing the ‘aether’ in terms of the motions of particles), and the eventual abandonment of mechanical models and reliance on field concepts as fundamental.”⁹⁴ So perhaps many of the issues that I have described are really based on an inaccurate conception of physics to begin with. Who knows. In any case, I believe that the difficulties that seem to arise as soon as we try to give a mechanical interpretation of vision, along the more familiar reductionist lines that I have proposed, and that seems the standard in most of the philosophical literature, are certainly instructive.

V

It has been my concern to show that vision is straightforwardly a process of representation. By allowing the distant environment to come into view, vision allows the organism to anticipate things and events that, in a sense, are still absent. This is essentially what makes vision a process of representation. Although this interpretation of the representational nature of vision is somewhat unorthodox, in that it does neither involve, or is dependent on, the claim that vision is indirect nor the claim that it is intrinsically contentful, the problems that we are left with are still recognizably the kind of problems

⁹⁴See e.g. (Bilodeau; 1996, p. 387).

that we expect instances of representation to give rise to. Primarily the interaction problem that is posed by the action at a distance that is suggested by the phenomenology of vision is problematic in this respect, and indeed much reminiscent of the more traditional problem of the causal efficacy of content. I have tried to show why I think that the phenomenology of vision is difficult to accommodate, given what we suppose the physical world is like. The apparent action at a distance that is at work in vision appears to straightforwardly contradict the familiar naturalistic assumption that the only direct causal determinants of human thought and behavior are spatiotemporally local. I have tried to give some reasons to suppose that eliminativism and reductionism may be inadequate in the face of this contradiction. Alternative metaphysical theories that suggest themselves are emergentism and holism. Alternatively, developments in physics might suggest that mechanism and locality are bad characterizations of the physical to begin with. Whatever it might be, we need to come to grips with, and acknowledge, the apparent action at distance that is suggested by visual phenomenology and that renders vision as essentially a process of representation.

References

- Aydede, M. (2014). Pain, *Stanford Encyclopedia of Philosophy*.
URL: <http://plato.stanford.edu/entries/pain/>
- Barsalou, L. W. (1999). Perceptual symbol systems, *Behavioral and brain sciences* **22**(04): 577–660.
- Berkovitz, J. (2014). Action at a distance in quantum mechanics.
URL: <http://plato.stanford.edu/archives/spr2014/entries/qm-action-distance/>
- Bilodeau, D. J. (1996). Physics, machines, and the hard problem, *Journal of Consciousness studies* **3**(5-6): 386–401.

- Block, N. (2003). Mental paint, *Reflections and replies: Essays on the philosophy of Tyler Burge* .
- Block, N. (2006). Bodily sensations as an obstacle to representationism, in M. Aydede (ed.), *Pain: New Essays on Its Nature and the Methodology of Its Study*, MIT Press, Cambridge, MA, chapter 6.
- Brewer, B. (2006). Perception and content, *European Journal of Philosophy* **14**(2): 165–181.
- Broad, C. D. (1949). The relevance of psychical research to philosophy, *Philosophy* **24**(91): 291–309.
- Brooks, R. A. (1991). Intelligence without representation, *Artificial intelligence* **47**(1): 139–159.
- Campbell, J. (2002). *Reference and Consciousness*, Oxford University Press, Oxford, UK.
- Chalmers, D. (1996). *The Conscious Mind: In Search of a Fundamental Theory*, Oxford University Press, Oxford, UK.
- Chalmers, D. (1998). Facing up to the problem of consciousness, *Consciousness and Emotion in Cognitive Science: Conceptual and Empirical Issues* **3**: 207.
- Chalmers, D. (2006). Strong and weak emergence, in P. D. . P. Clayton (ed.), *The reemergence of emergence*, Oxford University Press, New York, NY, pp. 244–256.
- Chemero, A. (2000). Anti-representationalism and the dynamical stance, *Philosophy of Science* pp. 625–647.
- Chemero, C. (2009). *Radical Embodied Cognitive Science*, MIT Press, Londond, UK.

- Churchland, P. M. (1985). Reduction, qualia, and the direct introspection of brain states, *The Journal of Philosophy* pp. 8–28.
- Churchland, P. M. (1989). *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*, MIT Press, Cambridge, MA.
- Clark, A. (1989). Microfunctionalism: Connectionism and the scientific explanation of mental states.
URL: <http://www.philosophy.ed.ac.uk/people/clark/pubs/microfx.pdf>
- Clark, A. and Chalmers, D. (1998). The extended mind, *analysis* pp. 7–19.
- Clark, A. and Toribio, J. (1994). Doing without representing?, *Synthese* **101**(3): 401–431.
- Colombo, M. (2014). Neural representationalism, the hard problem of content and vitiated verdicts. a reply to hutto & myin (2013), *Phenomenology and the Cognitive Sciences* **13**(2): 257–274.
- Copeland, J. (1996). What is computation?, *Synthese* **108**(3): 335–359.
- Crane, T. (1995). *The Mechanical Mind: A Philosophical Introduction to Minds, Machines and Mental Representation*, Routledge, London, UK.
- Degenaar, J. and Myin, E. (2014). Representation-hunger reconsidered, *Synthese* **191**(15): 1–10.
- Dennett, D. (1987). *The Intentional Stance*, The MIT Press, Cambridge, MA.
- Dennett, D. (1988). Quining qualia, in A. Marcel and E. Bisiach (eds), *Consciousness in Contemporary Science*, Oxford University Press, New York, Ny, chapter 17.
- Dennett, D. (1991). *Consciousness Explained*, Little, Brown and Company, Boston, MA.

- Dennett, D. (1998). *Brainchildren: Essays on designing minds*, MIT Press, Cambridge, Ma.
- Dennett, D. and Haugeland, J. (1987). Intentionality.
URL: <http://ase.tufts.edu/cogstud/dennett/papers/intentio.htm>
- Dretske, F. (1995). *Naturalizing the Mind*, MIT Press, London, Eng.
- Evans, G. (1982). *The varieties of reference*, Clarendon Press Oxford, New York, NY.
- Fish, W. (2010). *Philosophy of perception: A contemporary introduction*, Routledge, New York, NY.
- Fodor, J. (1975). *The language of thought*, Thomas Y Crowell Company, New York, NY.
- Fodor, J. (1981). *Representations: Philosophical essays on the foundations of cognitive science*, The Harvest Press Limited, Brighton, UK.
- Fodor, J. A. (1987). *Psychosemantics: The problem of meaning in the philosophy of mind*, MIT Press, Cambridge, MA.
- Fodor, J. A. (1994). *The Elm and the Expert*, MIT Press, Cambridge, Ma.
- Fresco, N. (2010). Explaining computation without semantics: keeping it simple, *Minds and Machines* **20**(2): 165–181.
- Gibson, J. J. (1972). A theory of direct visual perception, in A. Noë and E. Thompson (eds), *Vision and Mind: Selected Readings in the Philosophy of Perception*, MIT Press, Cambridge, MA, chapter 5.
- Griffin, D. R. (1993). Parapsychology and philosophy: A whiteheadian post-modern perspective, *Journal of the American Society for Psychical Research* **87**: 217–88.

- Güzeldere, G. (1997). The many faces of consciousness: A field guide, *The nature of consciousness: Philosophical debates*, MIT press, Cambridge, MA, chapter Introduction.
- Harman, G. (1990). The intrinsic quality of experience, *Philosophical perspectives* **4**: 31–52.
- Harnad, S. (1990). The symbol grounding problem, *Physica D: Nonlinear Phenomena* **42**(1): 335–346.
- Haselager, P., de Groot, A. and van Rappard, H. (2003). Representationalism vs. anti-representationalism: a debate for the sake of appearance, *Philosophical psychology* **16**(1): 5–24.
- Heidegger, M. (1967). *Sein und Zeit*, 11th edn, Max Niemeyer Verlag, Tübingen, GER.
- Hill, C. S. (2009). *Consciousness*, Cambridge University Press, New York, NY.
- 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*, Oxford University Press, chapter 49.
- Horst, S. (1999). Symbols and computation a critique of the computational theory of mind, *Minds and Machines* **9**(3): 347–381.
- Hutto, D. D. (1999). *The Presence of Mind*, John Benjamins Publishing, Amsterdam, NL.
- Hutto, D. D. (2011). Enactivism: Why be radical?, *Sehen und Handeln* pp. 21–44.
- Hutto, D. D. (2013). Why believe in contentless beliefs?, in N. Nottelmann (ed.), *New Essays on Belief: Constitution, Content and Structure*, Palgrave Macmillan, London, UK, chapter 3.

- Hutto, D. D. and Myin, E. (2013). *Radicalizing enactivism: Basic minds without content*, MIT Press, Cambridge, MA.
- Jackson, F. (1977). *Perception: A representative theory*, CUP Archive, New York, NY.
- Jackson, F. (1982). Epiphenomenal qualia, *The Philosophical Quarterly* **32**(April): 127–136.
- Kennedy, M. (2009). Heirs of nothing: The implications of transparency, *Philosophy and Phenomenological Research* **79**(3): 574–604.
- Kim, J. (2000). *Mind in a physical world: An essay on the mind-body problem and mental causation*, MIT press, Cambridge, MA.
- Kim, J. (2006). Emergence: core ideas and issues, *Synthese* **151**(3): 547–559.
- Lehar, S. (2015). Corrupted definition of representationalism by tye and dretske.
URL: <http://cns-alumni.bu.edu/~slehar/epist/corruption.html>
- McCulloch, W. S. and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity, *The bulletin of mathematical biophysics* **5**(4): 115–133.
- McDowell, J. (1986). Singular thought and the extent of inner space, in J. McDowell and P. Pettit (eds), *Subject, Thought, and Context*, Oxford University Press, New York, NY, chapter 5.
- McDowell, J. (1994). The content of perceptual experience, *The Philosophical Quarterly* pp. 190–205.
- McDowell, J. H. (1998). *Mind, value, and reality*, Harvard University Press, Cambridge, MA.
- Mendelovici, A. and Bourget, D. (2014). Naturalizing intentionality: Tracking theories versus phenomenal intentionality theories, *Philosophy Compass* **9**(5): 325–337.

- Newell, A. and Simon, H. A. (1976). Computer science as empirical inquiry: Symbols and search, *Communications of the ACM* **19**(3): 113–126.
- Noë, A. (2006). Experience without the head, in T. S. Gendler and H. John (eds), *Perceptual Experience*, Oxford University Press, New York, NY, chapter 11.
- Orlandi, N. (2014). *The Innocent Eye: Why Vision is Not a Cognitive Process*, Oxford University Press, New York, NY.
- Peacocke, C. (1983). *Sense and content: Experience, thought, and their relations*, Clarendon Press, Oxford, UK.
- Piccinini, G. (2008). Computation without representation, *Philosophical Studies* **137**(2): 205–241.
- Piccinini, G. (2009). Computationalism in the philosophy of mind, *Philosophy Compass* **4**(3): 515–532.
- Piccinini, G. (2012). Computation in physical systems.
URL: <http://plato.stanford.edu/archives/fall2012/entries/computation-physicalsystems/>
- Pinker, S. (1997). *How the mind works*, The Penguin Press, London, UK.
- Poggio, T. (1981). Marr’s computational approach to vision, *Trends in neurosciences* **4**: 258–262.
- Ramsey, W. M. (2007). *Representation reconsidered*, Cambridge University Press, Cambridge, UK.
- Rey, G. (1983). A reason for doubting the existence of consciousness, *Consciousness and self-regulation*, Springer, pp. 1–39.
- Robinson, H. (1994). *Perception*, Routledge, London, UK.
- Robinson, H. (2008). Why frank should not have jilted mary, in E. L. Wright (ed.), *The Case for Qualia*, MIT Press, chapter 12.

- Shapiro, L. (2011). *Embodied Cognition*, Routledge, New York, UK.
- Shear, J. (1997). *Explaining consciousness*, MIT Press, Cambridge, MA.
- Shoemaker, S. (2003). Content, character and color, *Philosophical Issues* **13**(1): 253–278.
- Smolensky, P. (1988). On the proper treatment of connectionism, *Behavioral and brain sciences* **11**(01): 1–23.
- Snowdon, P. (1980). Perception, vision and causation, in A. Noë and E. Thompson (eds), *Vision and Mind: Selected Readings in the Philosophy of Perception*, MIT Press, Cambridge, MA, chapter 9.
- Stich, S. (1992). What is a theory of mental representation?, *Mind* **101**(402): 243–261.
- Stich, S. P. (1983). *From folk psychology to cognitive science: The case against belief.*, the MIT press, Cambridge, MA.
- Stoljar, D. (2010). *Physicalism*, Routledge, New York, NY.
- Thompson, B. (2008). Representationalism and the argument from hallucination, *Pacific Philosophical Quarterly* **89**(3): 384–412.
- Travis, C. (2004). The silence of the senses, *Mind* **113**(449): 57–94.
- Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*, MIT Press, London, UK.
- Tye, M. (2002). Representationalism and the transparency of experience, *Noûs* **36**(1): 137–151.
- Tye, M. (2014). Transparency, qualia realism and representationalism, *Philosophical Studies* **170**(1): 39–57.
- Van Gelder, T. (1995). What might cognition be, if not computation?, *The Journal of Philosophy* **92**(7): 345–381.

van Rooij, I., Bongers, R. M. and Haselager, W. P. F. (2002). A non-representational approach to imagined action, *Cognitive Science* **26**(3): 345–375.

Wittgenstein, L. (1989). *Tractatus logico-philosophicus*, 5th edn, Athenaeum – Polak & Van Genneep, Amsterdam, NL.

Wright, E. L. (2008). *The Case for Qualia*, MIT Press, Cambridge, MA.