

The Perception/Cognition Divide: *One More Time, with Feeling*

Uriah Kriegel

Abstract. Traditional accounts of the perception/cognition divide tend to draw it in terms of subpersonal psychological processes, processes into which the subject has no first-person insight. Whatever betides such accounts, there seems to *also* be some first-personally accessible difference between perception and thought. At least in normal circumstances, naïve subjects can typically tell apart their perceptual states from their cognitive or intellectual ones. What are such subjects picking up on when they do so? This paper is an inconclusive search for an answer. At its end, I conclude, without joy, that we may have to simply accept the perception/cognition distinction as a primitive and inexplicable bright line within the sphere of conscious phenomena.

1 *Introduction: The Perception/Cognition Divide from the First-Person Perspective*

Suppose you are in a conscious mental state that represents a brown dog, and an omniscient authority assures you that your mental state is either a *perceptual experience* of a brown dog or an *occurrent thought* about a brown dog – it is not a memory of a brown dog, not an imagination of a brown dog, not a fear of a brown dog, and so on. Then the authority asks you whether, in your opinion, your mental state is in fact a perceptual experience or a thought.

There are two assumptions about this scenario that I would like to make from the outset. The first is that you should be able to answer the authority's question fairly effortlessly and immediately, with a considerable degree of reliability. Perhaps you will not be infallible in answering such questions, and perhaps a sufficiently inventive evil demon or scientist could put you in sufficiently atypical circumstances that you would get it wrong often. But in typical circumstances, it seems most of us will, most of the time, be right about whether our brown-dog-representing state is a perceptual experience or a thought.

The second assumption I want to make is that you can answer the authority's question without access to any third-person evidence. For present purposes, this includes four types of evidence: verbal report, nonverbal behavior, third-party testimony, and brain scans. Typically, when asked whether *someone else* is perceiving a brown dog or thinking about a brown dog, you need access to at least one of these four types of information to take any informed view; but when asked whether *your own* mental state is a perceptual experience or a thought, you can have justified belief on the matter even in the absence of evidence from any of these four groups.

I do not think these two assumptions are beyond reasonable doubt, and will in fact take a more critical approach to them in the final section of this paper. But the two assumptions are on their faces rather plausible. And if we accept them both, we face the following question: What feature(s) of your mental state do you pick up on when you answer the authority's question? If we can answer such questions relatively effortlessly and reliably, despite having no *third-person* evidence of relevance, it would seem that we must have *first-person* insight into the perceptual or cognitive character of our conscious states. But what is the exact content of this first-person insight?

In the 1970s and 80s, philosophers of mind developed a number of sophisticated accounts of the perception/cognition divide. However, these accounts tended to draw the distinction in terms of subpersonal or architectural features of mental states and processes, ones into which naïve subjects do not seem to have any first-person insight. Consider the Fodorian notion that perceptual processes are 'modular' whereas cognitive ones are 'central.' Fodor (1983) adduced a number of features distinguishing the former from the latter, notably domain-specificity and informational encapsulation. However, these do not seem to be features into which we

have first-person insight (nor did Fodor intend them to be such). Domain-specificity is a matter of responsiveness to a limited range of inputs, hence a matter of the origin of the processes producing the relevant mental states; informational encapsulation is a matter of susceptibility to cognitive penetration, that is, to influence from certain top-down processes. Clearly, an ordinary subject may have no pertinent information about the processes in which her current mental representation of a brown dog originated, or the processes to which it is susceptible, and yet have the kind of first-person insight into certain features of her current mental state that makes it possible for her to reliably tell whether that state is perceptual or cognitive.

In what follows, I will call such features *first-personally manifest* features. The question I want to raise, then, is this: What are the first-personally manifest features that distinguish perceptual experiences and occurrent thoughts? The architectural and subpersonal approaches from a generation ago may not answer this question (or if they do, we would still need to get clear on what the answer *is*). What might?

2 *The Perception/Cognition Divide and Intentional Character*

One straightforward approach to our problem might be: perceptual experiences have one kind of phenomenal character, thoughts have another, and this is what a naïve subject picks up on when forming a justified belief on the matter. Ultimately, this answer may well be right. But absent further elaboration, all it says is this: perception has a perceptual phenomenal character, cognition has a cognitive phenomenal character. This does not yet illuminate the nature of the perceptual/cognitive divide.

Indeed, this helps us see our problem from another angle. In the 70s and 80s, the debate on the perception/cognition divide was conducted against the background of a widespread assumption that while perceptual experiences have phenomenal character, thoughts are non-phenomenal states. Since those days, many philosophers have come to countenance a proprietary kind of cognitive phenomenology, that is, a *sui generis* phenomenal character proper to thought (see, e.g., Siewert 1998, Pitt 2004). But once we recognize cognitive phenomenology alongside perceptual

phenomenology, we incur the theoretical debt of accounting for the phenomenal *difference* between the two. So, even if the features our subject picks up on, when classifying her occurrent mental representations as perceptual or cognitive, are certain phenomenal properties, there is still a legitimate kind of curiosity we might have about the *nature* of these phenomenal properties – and of the difference between them. Suppose all and only perceptual experiences instantiate phenomenal property P, while all and only occurrent thoughts instantiate phenomenal property T. It is natural to want to know *what P and T are*. On the assumption that phenomenology is by its nature first-personally manifest, providing an account of the difference between P and T – between perceptual and cognitive phenomenology – is also answering the question of what is the first-personally manifest difference between perception and thought.

It is also possible, of course, to *deny* the existence of any proprietary cognitive phenomenology. If one does, one instantly acquires a straightforward answer to our question: perceptual experiences have phenomenal character, thoughts do not – that is the first-personally manifest difference between them. I will be assuming here, however, that thoughts do have a phenomenal character proper to them. I concede that it is only against the background of this assumption that our question becomes pressing.

Now, at one level, I think the phenomenal feature that distinguishes a perceptual experience of a brown dog from a thought about a brown dog is pretty straightforward. The problem is that the most natural ways of capturing this feature are metaphorical. We say, for example, that in perceptual experience objects are *directly* present to us, whereas in thought they are not. But what does ‘directly’ exactly *mean*? Husserl puts it nicely when he says that while perceptual experiences present their objects *in the flesh* (*‘in persona,’* he writes), thoughts do not. But that is just a compelling metaphor. Our question is how to get underneath such suggestive expressions and provide a substantive, informative account of the first-personally manifest difference between perception and thought.

There are many ‘leads’ one might consider. One might, for example, be visited by a hunch that perception is crucially ‘perspectival’ whereas thought is ‘objective’; or one might feel that the key difference is that perception is in some sense passive whereas thought is active (*‘receptivity’* vs. *‘spontaneity’*); or one might be impressed by

the fact that perception seems to be phenomenally rich whereas thought is phenomenally sparse at best. I cannot hope to follow up here on *all* the possible leads of relevance. To restrict the group of options to be considered in a principled way, I will focus on features of perceptual experiences and occurrent thoughts that pertain to their *intentional character*. By this I mean both (i) these states' *intentional content* and (ii) the kinds of *attitude* they take toward these contents. (There are different views on the relationship between content and attitude. One model treats contents and attitudes as distinct features that combine together to form intentional characters; another model takes the basic intentional structure to be that of taking-attitude-A-toward-content-C and treats content and attitude as theoretical abstractions of some sort. Although I lean toward the second understanding, I will take no official stand on this question here. What matters to me is that not only content but also attitude be examined when we consider the intentional profile of mental states.)

The restriction to intentional character is not *too* restrictive. Many philosophers today are impressed by the idea of the 'transparency of experience' (Harman 1990). As is well known, though, some very different claims have been associated with this label. Most strongly, some hold that when we first-personally attend to our own conscious states, we are only aware of the external objects and/or properties these states are intentionally directed at. More weakly, some hold that what is first-personally manifest when we attend to our conscious states are properties of these states themselves, but only properties of the form having-such-and-such-intentional-content. More permissively yet, someone might hold that among the first-personally manifest properties of conscious states are not only content properties but also attitudinal properties. Indeed, this view is so permissive that many might consider it not to qualify as a transparency view at all. Thus in considering both content and attitude views we are imposing a relatively permissive restriction on the space of possible approaches.

It is possible, of course, that although perceptual experiences and occurrent thoughts are different in a first-personally manifest way, they are so in virtue of features that are neither content features nor attitudinal features. This possibility will not be ruled out by anything I will say hereafter – that is the cost of our restriction.

In what follows, I will devote more attention to attitudinal approaches to the first-personally manifest divide between perception and cognition, because elsewhere

(Kriegel forthcoming) I argue in some detail against content approaches. Nonetheless, I start, in §3, with a summary of my case against content approaches. In §4, I introduce and motivate the attitudinal approach. In §§5-6, I consider two specific attitudinal accounts, arguing against each. In §7, I consider our options given the failure of both the content and attitudinal approaches examined here; I conclude, without joy, that we may be forced to adopt a primitivist position according to which there is simply two distinct first-personally manifest features, a *sui generis* phenomenology of perceptuality and a *sui generis* phenomenology of cognitivity, such that perceptual experiences exhibit the former whereas conscious thoughts exhibit the latter.

Before starting, three clarificatory comments are in order. First, I will use the expressions 'cognition' and 'thought' interchangeably, and have in mind such phenomena as thinking to oneself that the weather is nice, considering the likelihood that Spain will win the next world cup, and realizing that the deadline for submitting a grant proposal is only ten days away. There is a use of the term 'cognition' that makes it considerably wider, covering also phenomena of mental imagery, episodic memory, and any other stimulus-independent mental phenomenon (see Phillips forthcoming). Indeed, there is a use of 'cognitive' where it means essentially the same as 'mental' (see under: 'cognitive science'). I will not be using the term 'cognitive' in these more extended senses. For that would effectively turn the question before us into the more general one of what distinguishes perceptual from non-perceptual mental phenomena. But *that* question is really an amalgam of a number of separate questions: (a) what distinguishes perception from imagination?; (b) what distinguishes perception from episodic memory?; (c) what distinguishes perception from thought?; and so on. I think the right approach here is to take up each of these more focused questions in its own turn and give it the time of day. Running them together risks papering over important distinctions and producing a more confused picture of the main bright lines within the overall field of mental phenomena.

Secondly, it is important in some contexts to distinguish between three very different kinds of thought reports. The locution 'thinking *that*' tends to figure in reports of mental states that take a stand on the truth of that-which-is-thought-about, such as judging that *p* and believing that *p*. In contrast, the locution 'thinking *of*' tends to figure in reports of mental states that do *not* take a stand on the question of truth, such as entertaining *p* and contemplating *p*. Entertaining or contemplating *p* does not

constitutively involve taking p to be true, but judging or believing that p does. There is also a locution that tends to be neutral on the matter, namely, 'thinking *about*.' In any case, in what follows I will restrict myself to the think-that and think-about locutions, and have in mind only the committal, belief-like states, not the neutral, entertaining-like states. The reason for this is that I take perception to be likewise committal regarding the reality of its objects – more on this in §4. Therefore, it is for the contrast between perception and *committal* thought that the question of distinguishing the two presents itself most urgently.

Third comment: the discussion to follow ignores entirely the phenomena of *unconscious* perception and thought. I will not consider such phenomena as blindsight or even subliminal perception, nor tacit or dispositional belief. The reason for this is that I do not expect there to be any *first-personally manifest* difference between such mental states. For these, the sub-personal and architectural approaches from the 80s will do just fine. The residual problem I have identified at the opening does not arise for them. So although I will for the sake of brevity speak of perception and thought, I really just mean conscious perceptual experience and conscious occurrent thought ('cognitive experience'). (Note: for the sake of flow I use 'perception' to cover veridical, illusory, and hallucinatory perceptual experiences alike; that is, I bracket the fact that the term is, in its most central usage, a success term.)

3 *The Content Strategy*

There is a natural temptation to distinguish perception and thought in terms of what they represent. In particular, we might think, perception characteristically represents *low-level* properties, such as color and shape, whereas thought typically represents *high-level* properties, such as being a dog and being Parisian. More generally, the idea is to start by dividing represented entities into *sensibles* and *intelligibles*, then explain the difference between perception and thought in terms of representing the former versus representing the latter.

The main problem here is that while there may be restrictions on what perception can represent, there appear to be no restrictions on what thought can represent. In particular, it is unclear why, whatever the property one's perception represents, one could not simply think about the-property-one's-perception-represents, if only under that description. In that scenario, there might be a difference in the *description used to pick out* this property, but there will be no difference in the property being picked out.

Note that it is beside the point that there is still a difference here insofar as certain things perception cannot represent whereas thought can represent anything representable. For that only helps us to distinguish the faculties of thought and perception, but does not help us distinguish two token mental states, one perceptual and one cognitive, which both represent F. Yet a normal subject in normal conditions can tell whether her mental state is a perceptual experience of red or a thought about red.

The reasoning so far suggests that the content strategy would do well to draw the perception/cognition distinction not in terms of what properties are being represented, but in terms of something like the *modes of presentation* used to represent those properties. Of course, it would be uninformative to simply state that perception employs a perceptual mode of presentation whereas thought employs a cognitive mode of presentation. It would likewise be of limited help to say that perception employs an in-the-flesh mode of presentation whereas thought does not. What we want is to capture the psychological reality underlying that metaphor. It is more informative to unpack mode of presentation either in terms of abstract entities in Frege's 'third realm' or in terms of functional role. But both of these do not seem like first-personally manifest features. Certainly abstract objects cannot be *introspected*, since introspection directs itself at one's *current* stream of consciousness, not at a realm of timeless beings. And arguably, functional role cannot be introspected either, since functional role is a *dispositional* property, whereas introspective awareness presents only categorical properties. (Compare: you cannot see the vase's fragility, even if in some sense you can see *that* the vase is fragile. Likewise, I would argue, you cannot introspect your experience's disposition to elicit such-and-such behaviors and inferences, even if you can introspect *that* your experience has such a disposition.)

It is hard to see, then, how to produce an informative-yet-adequate characterization of the alleged difference between the modes of presentation employed by perception and by thought. Might one suggest that perception employs *perspectival* modes of presentation, whereas thought employs non-perspectival ones? This is problematic without further explanation of what makes a mode *perspectival*. If we understand it, as is common, in terms of *indexical representation*, then the resulting account would be extensionally inadequate: many thoughts have indexical content as well. In addition, it is unclear that indexicality and *perspectivity* characterize *all* perception. For example, an olfactory experience of the odor of freshly ground coffee (without a clear sense of where the odor is coming from, say) is certainly perceptual, but in what sense is it *perspectival*? We are still left with the question of what distinguishes that experience from a conscious occurrent thought about the odor of freshly ground coffee.

It might be suggested that the key difference between perceptual and cognitive content is more structural: thought employs *conceptual* modes of presentation, perception *nonconceptual* ones. The notion that perception has *nonconceptual* content whereas thought has *conceptual* content is a recurring theme in contemporary philosophy of mind (e.g., Dretske 1995). There are several different ways this notion can play out (Heck 2000), but the main one is probably this: a perceptual experience of a brown dog is a representational state we can be in even when we do not possess the concepts BROWN and DOG, whereas an occurrent thought about a brown dog is a representational state that we cannot be in unless we possess BROWN and DOG.

An immediate concern is that the suggestion offers only a *negative* characterization of perceptual experience: it is a certain kind of lack, or absence, that makes a mental state perceptual. This seems at best incomplete, as there also seems to be something positive *present* in perceptual experience – all those vividly present colors and sounds we experience! Arguably, it is precisely this ‘something positive’ that a naïve subject picks up on when classifying her conscious state as perceptual. One could of course use the label ‘nonconceptual’ for something positive, but one would also have to make sure that that positive feature is first-personally manifest. Stalnaker (1998), for instance, understands nonconceptual content in terms of functions from worlds to extensions; but surely we cannot *introspect* such functions, so it is unclear in what sense they might be first-personally manifest.

In fact, it is questionable in general that the conceptual/nonconceptual distinction is first-personally manifest. On many accounts, it is ultimately a difference in *functional role*: a conceptual representation of a hexagon implicates certain capacities, such as to recognize and reidentify hexagons across time, as well as trigger certain inferences that nonconceptual representations of hexagons do not. But to repeat, functional role, as a dispositional property, is not an introspectible feature.

The biggest problem with the account under consideration, however, is that the nonconceptual character of our perceptual experiences appears to be a merely contingent and accidental feature of theirs. Imagine a supersentient creature – call it Lynceus – who possesses a concept for every shade of red and every polygon. Lynceus can discriminate red₂₇₃ from red₂₇₄, recognize a red₂₇₃ sample a year after seeing it, and draw appropriate inferences about objects' color properties. Imagine now that Lynceus is presented with a red₂₇₃ megagon (in good lighting conditions etc.). It would be strange to say that Lynceus cannot perceptually experience the red₂₇₃ megagon, but can only *think* about it, on the grounds that his perceptual acuity and processing power is so much better than ours. Perceptual acuity helps us see more, not less! But if Lynceus possesses the concepts for the properties he perceives, then it is not *in the nature* of perception to outstrip the perceiving subject's conceptual capacities. Crucially, there is no reason to expect that someone like Lynceus be unable to distinguish his perceptual experiences of red₂₇₃ megagons from his thoughts thereabout. So he must pick up on some first-personally manifest difference between them that has nothing to do with whether concepts are used to pick out the properties of being red₂₇₃ and being a megagon. Upshot: it is a merely contingent fact about *us* that we fail to possess concepts for many features we can perceive, a fact that tells us nothing about the nature of perception as such.

For the same reason, it is implausible to draw the perception/cognition distinction directly in terms of what often *motivates* the conceptual/nonconceptual approach, namely, the notion that perceptual experience is informationally rich whereas thought is informationally sparse. It is true that a typical perceptual experience of a yellow cornfield contains much detail that goes beyond yellowness and corniness, whereas a thought to the effect that the cornfield is yellow, or that there is a yellow cornfield about, does not. But again, this seems merely contingent: a Laplacean demon may have extremely detailed thoughts, while a perceptual experience of a yellow

Ganzfeld boasts a rather sparse content – yet there is obviously something different, manifest to the first-person perspective, between such states and their perceptual and cognitive counterparts. In addition, a richness difference would be a difference in degree, not a difference in kind; whereas the experiential difference between thought and perception seems to be rather a difference in kind. Perception does not seem to be cut of the same cloth as thought but just have much more of that cloth, so to speak. This calls for some *categorical* distinction between the two.

There are other possible content approaches to the perception/cognition divide. As noted, I provide a more indulging discussion of the approach in Kriegel Ms. Here I will rest content with what I have said so far against the content approach, moving now to consider the attitudinal approach.

4 *The Attitudinal Strategy: Background*

If the contents of perception and thought could be shown to always allow a perfect match, then the first-personally manifest difference between perceptual experience and thought would be proven not to be a content difference. But it would still not follow that it is not an intentional or representational difference. For it might be a difference at the level of *attitude*. So-called impure intentionalists (Chalmers 2004, Crane 2009) have held that some phenomenal features of conscious states pertain to attitude. Since phenomenal features are paradigmatically first-personally manifest, this has the potential to yield a satisfactory approach to our problem.

Obviously, construals of attitude in terms of functional role will raise the aforementioned problems concerning introspectibility. And such construals have certainly been prominent in philosophy of mind. When we construe the difference between belief that p and desire that p in terms of whether p is in the 'belief box' or the 'desire box,' the background picture is the kind of boxes-and-arrows psychology that characterizes mental phenomena in terms of their functional role within the overall economy of mental life.

But functional role is not the only way to construe attitudes. Thus, it is sometimes said that whereas belief that p represents p as *obtaining*, desire that p represents p rather as *good*, in some suitably generic sense of 'good.' Thus, Dennis Stampe writes:

[W]hile the belief and the desire that p have the same propositional content and represent the same state of affairs, there is a difference in the way it is represented in the two states of mind. In belief it is represented as *obtaining*, whereas in desire, it is represented as a state of affairs *the obtaining of which would be good*. This *modal* difference explains why a desire that p is a reason to make it true that p , while the belief that p is not. (Stampe 1987: 355; see also Tenenbaum 2009: 96)

Note that this way of describing attitudinal facts casts them as clearly *representational facts*: they are facts about belief's and desire's distinctive manners of representing their content. To be sure, the content of the belief that the weather will be nice in April is simply <the weather will be nice in April>, not <the state of affairs of the weather being nice in April obtains>. But this is precisely why the information about the state of affairs' obtaining is encoded into the very attitude of believing, not the belief's content. We may put this by saying that the belief *represents-as-obtaining* the content <the weather will be nice in April>. In the expression 'represents-as-obtaining,' the 'as obtaining' denotes a modification of the representation relation; it does not qualify that which is being represented. The desire that the weather will be nice in April, meanwhile, *represents-as-good* the same content. Representing-as-obtaining and representing-as-good are the distinctive attitudinal characters of belief and desire, on this approach, and they are *representational facts* about belief and desire – though not facts about the representational *contents* of belief and desire.

This way of thinking of 'attitudinal character' may open the door to an attitudinal account of the first-personally manifest difference between perceptual experiences and occurrent thoughts: given a perception and a thought that represent the same state of affairs (e.g., a rectangle's being red), the task is to identify the right P and T, such that the perception represent-as-P the rectangle's being red whereas the thought represents-as-T the rectangle's being red.

One general objection to this approach is that the nature of properties of the form representing-as-F is unclear. However, while this is a legitimate complaint, I will

set it aside here and argue that there are more specific problems with concrete attitudinal accounts of the perception/cognition divide.

We have already mentioned in passing that the attitudinal character of belief is probably that of representing-as-obtaining. Plausibly, belief that p shares its attitudinal character with such mental states as judging that p and thinking that p . For all of these, it is natural to assign representing-as-obtaining as attitudinal character.¹ If we could identify a *different* attitudinal character associated with perceptual experience, this would give us an initial account of the first-personally manifest perception/cognition divide.

One suggestion might be that perceiving a red rectangle is a matter of representing-as-existent, or representing-as-real, the red rectangle. Perhaps this would be claimed to explain, for example, the phenomenal contrast between perceptual experience of a red rectangle and (sensory) imagination of a qualitatively indistinguishable rectangle: while the former represents-as-real the rectangle, the latter represents-as-unreal the rectangle (Kriegel 2015a). If this is right, then the cognition/perception divide is to be drawn in terms of representing-as-obtaining vs. representing-as-real/existent.

But this account of the perceptual attitude is doubly problematic. While representing-as-existent is suitable for perception of *concrete particulars*, but arguably, there is also such a thing as perception of *states of affairs*. Looking in the relevant direction, you see not only the rectangle but also its redness, indeed you see the redness as being *the rectangle's* redness. Presumably, perception of states of affairs would involve representing-as-obtaining rather than representing-as-existent, insofar as the mode of being of states of affairs is that of obtaining rather than existing. To that extent, the proposed suggestion could only work against the background of the assumption that states of affairs cannot be perceived. But this is an odd commitment to take on board from the outset. (Note well: if we think that the difference in mode of being between concrete particulars and states of affairs is terminological and insubstantive, then the proposed account would itself become insubstantive, since what it proposes is distinguishing thought and perception in terms of representing-as-obtaining vs. representing-as-existent.)

The general point is: for at least some thought about a 's being F (e.g., some rectangle's being red), there is a corresponding perceptual experience of a 's being F . In that scenario, the two mental states represent the very same entity, indeed both represent-*as-obtaining* that entity; yet there is intuitively a stark difference between seeing and thinking of some rectangle's being red, one that should be introspectively manifest to any subject capable of having either state.

5 Spatiotemporal Attitudinal Characters

A mental state's attitudinal character goes with a specific kind of correctness conditions. If p obtains but is bad, for example, then it is correct to represent-*as-obtaining* p but not to represent-*as-good* p , hence to believe but not to desire that p . If instead p is good but fails to obtain, then it is correct to desire but not to believe that p . This creates an interesting problem for ascribing to perceptual experiences the attitudinal character of representing-*as-obtaining*. Imagine you are hallucinating a rectangle, and hallucinate its being red. Imagine further, however, that, as it happens, there really is an intrinsically indistinguishable red rectangle, but on Mars rather than in front of you; or perhaps there really was such a rectangle exactly where you hallucinate it to be, but a year ago rather than now. Intuitively, your perceptual experience is *incorrect*. That is, its correctness conditions are not met. But if the experience's attitudinal character were that of representing-*as-obtaining*, then given that the Martian rectangle *exists*, and the state of affairs of its being red obtains (as does, at least by eternalist lights, the yesteryear state of affairs of the rectangle's being red), the experience would be correct. Thus the proposed account returns the wrong result that your hallucination is correct. This is what we might call *thought-experimental disconfirmation* of the account.

This problem may inspire the following suggestion: the difference between a perceptual experience of and a thought about some rectangle's being red is that the perceptual experience represents-*as-obtaining-here-and-now* the rectangle's being red, whereas the thought represents-*as-obtaining-simpliciter* the rectangle's being red. Here the distinction is between two ways of representing-*as-obtaining* a state of affairs: an unrestrictive way characteristic of thought, and a spatiotemporally restrictive way we

find in perception. We might say that perception exhibits an *egocentric* form of representing-as-obtaining, whereas thought exhibits an *allocentric* form of representing-as-obtaining.

The advantage of this account is that it neutralizes the Martian and yesteryear red rectangles, and allows the perception of states of affairs. Having a perceptual experience as of a rectangle's being red is incorrect, on the account, if no rectangle is red *here and now*. If a rectangle is red on Mars, but not here, then the experience's correctness conditions are not met; if a rectangle was red here, but a year ago rather than now, then again they are not met. Meanwhile, representing-as-obtaining-here-and-now is a kind of representing-as-obtaining, so can target states of affairs unproblematically.

For this account to enjoy any plausibility, however, it is important *not* to understand 'represents-as-obtaining-here-and-now' as denoting a species, or determinate, and 'represents-as-obtaining-*simpliciter*' as denoting the corresponding genus or determinable. For perception is not a species or determinate of thought. Accordingly, we should interpret the expression 'represents-as-obtaining-*simpliciter*' rather as meaning something like 'represents-as-obtaining-*somewhere-and-somewhen*'; and consider that both representing-as-obtaining-here-and-now and representing-as-obtaining-somewhere-and-somewhen are species of representing-as-obtaining. So understood, perception and thought can be seen to employ *incompatible* attitudinal characters – an egocentric one and an allocentric one.

To increase the account's plausibility, we might also consider weakening its characterization of the attitudinal character distinctive of perception as follows: instead of claiming that it is the character of representing-as-obtaining-here-and-now, one would put forward the disjunctive thesis that it is either representing-as-*obtaining*-here-and-now or representing-as-*existing*-here-and-now. The point of this is to allow perceptual experiences not only of states of affairs but also of individual objects. (Obviously, if we have positive reasons to think that perception of individual objects is impossible, at least outside the context of their being constituents of states of affairs, then this move is unnecessary.)

Even with these provisos in place, however, the 'here-and-now' account faces a number of difficulties. Note, to start, that the account does nothing to rule out the possibility of *thoughts* about the here and now. Fortunately: we can clearly *think* that there is a brown dog *here and now*. It is just that the here-and-now information conveyed by such a thought must be conveyed as part of its *content*. That is, while the perception of a rectangle's being red represents-as-obtaining-here-and-now the state of affairs of the rectangle's being red, the thought that the rectangle is red here and now represents-as-obtaining-*simpliciter* the (different) state of affairs of the rectangle's being red here and now. The first problem with the account under consideration is that it is far from clear that the difference between *content-based* here-and-now information and *attitudinally encoded* here-and-now information is first-personally manifest. Introspection alone does not seem to reveal whether a mental state represents x as F or represents-as-F x.

Another problem with the account under consideration is that it seems to get the extension wrong in certain cases. Someone may sit in a sports bar in Tokyo and watch a live broadcast of a basketball game taking place in Chicago. It is natural to say that this person is *seeing the game* – even though the game is not 'here' relative to this Tokyo-bound individual. Likewise, a person who looks at Mars through a telescope is naturally described as *seeing Mars*. But Mars is not here, unless we are so generous about what counts as here that we risk eroding the very distinction between 'here' and 'somewhere.'

These considerations raise a deeper problem for the here-and-now account: the account owes us some elucidation of what 'here' and 'now' actually *mean*, and while it is natural to elucidate 'now' as meaning something like 'simultaneous with the perceptual experience,' it is less clear what 'here' means. We can say that it means 'compresent with the experience,' or 'collocated with the experience,' but this is a vague and relative notion. We can imagine a superman with 'bionic vision' who, sitting among us, can see specific pebbles on the moon. What would count as 'here' for him would presumably be very different from what would count as 'here' for us. And this raises the specter of an even more extraordinary superwoman, one who can see to the end of the cosmos. Although the cosmos has been expanding very rapidly during the upward fifteen billion years since the Big Bang, the distances involved are still small potatoes for this superwoman. Such is the extraordinary strength of her vision, that only

in another five billion years she would start losing the ability to see certain faraway objects! (Note: it is of no consequence to this reasoning if this superwoman is nomologically impossible – so long as she is metaphysically possible.) It seems that even for our superwoman there is a difference between *thinking about* a specific grain of dust on Jupiter and *seeing* that grain of dust. And yet there is no part of the cosmos that is so far away from her that it could no longer be said to be ‘here’ relative to her – since all parts of the cosmos are *visible* (as well as audible etc.) to her.

In response, it might be suggested that we simply drop the ‘here’ part of the account, adopting the view that perception’s distinctive attitudinal character is simply that of representing-as-now. Given an understanding of ‘now’ in terms of simultaneity, this would amount to the following view: a perceptual experience represents-as-obtaining-simultaneously-with-its-own-occurrence the relevant state of affairs. Observe that this makes perceptual experience (but not thought) *token-reflexive*: each token perceptual experience figures in its own correctness conditions. However, the token-reflexivity here is due to perception’s attitude rather than content (compare Kriegel 2015b), and is a temporal rather than causal kind of token-reflexivity (contrast Searle 1983 Ch.2).

Unfortunately, there are many problems with this account. First, it brings back to life the problem of the Martian rectangle’s being red, which obtains now but fails to render your hallucinations of a rectangle’s being red correct or appropriate.

Secondly, there may be temporal analogs of the spatial problems generated by live broadcast and telescopic perception. A person may miss a concert due to a traffic jam, but be fortunate enough to still hear the concert, thanks to a recording played on the radio the next day. Another person, watching the sunset over the Mediterranean, sees the sun where and how it was almost eight minutes earlier (given that the sun lies almost eight light minutes away from Earth).

Most importantly, perhaps, the difference between representing-as-obtaining-now and representing-as-obtaining-*simpliciter* just does not seem like the right difference for capturing the distinctively ‘in the flesh’ character of perception. All those vivid colors and sounds directly present to the perceiving mind, but not to the thinking intellect, do not seem to have much to do with the simultaneity between perceptual

experience and its objects. (I am using 'object' here widely, to cover any type of entity that may be targeted by a mental state, not only concrete particulars.)

A final objection, which applies to both the hic-et-nunc account and the nunc-only account, has to do with so-called veridical hallucinations. Suppose that, under the influence of effective hallucinogens, you have a visual experience as of a butterfly fluttering exactly one meter straight ahead; and that while you are having this hallucination, a butterfly just like the one you are hallucinating enters the room and starts fluttering just where, and how, your hallucinated butterfly 'does' (Searle 1983 Ch.2). This seems like an incorrect visual experience. However, both the hic-et-nunc and nunc-only accounts return the result that the experience is a correct one. For the experience represents-as-obtaining(-here-and)-now the state of affairs of such-and-such a butterfly fluttering so-and-thus, and indeed such-and-such a butterfly is fluttering so-and-thus here and now.

6 *Causal Attitudinal Character*

On the face of it, the problem with the real butterfly before you is that it is in no way responsible for the experience you are having. For a perceptual experience to be correct, this suggests, it must be *caused by* the perceived object. Accordingly, we might opt for an account of perception that focuses not on a here-and-now condition on perceived objects, but on a *causal responsibility* condition – still understood as a condition imposed by the attitude, not content (see Recanati 2007 Ch.17). On this view, while the thought that the rectangle is red represents-as-obtaining the rectangle's being red, the corresponding perception represents-as-obtaining-and-causing-that-very-perception the rectangle's being red. On the plausible assumption that a state of affairs cannot cause something without obtaining, we might just say: the perception represents-as-causing-that-very-perception. Again we have a token-reflexive picture of perception (but not thought) here, and again this token-reflexivity is attitudinally grounded. It is just that here the token-reflexivity is causal rather than (spatio)temporal.

This causal-attitudinal account has several advantages. First, the notion of causality invoked here is easier to pin down than the notion of compresence or collocation suggested by 'here.'² Secondly, the causal account captures the pre-theoretically appealing idea that perception is stimulus-dependent in a way thought is not (Beck 2018, Phillips forthcoming), but instead of involving the *actual* stimulus-dependence of perceptual states – something which is presumably not *introspectively* manifest – it adduces an attitudinally imposed stimulus-dependence condition on perceptual objects. Thirdly, and most importantly, the causal account handles elegantly the cases of telescopic and live-broadcast perceptions (as well as the recording and sunset cases): in all these cases, although the perceived object is not in any useful sense here (or now), it is still causally responsible for the occurrence of the perceptual experience. For example, Mars is *causally responsible* for our stargazer's experience as she looks through the telescope; the basketball game in Chicago causes the broadcast, which in turn causes the Tokyo barfly's experience as she watches TV (and causation is transitive).

Of all the accounts we have considered in this paper, this one appears to me the most defensible. Nonetheless, there are three considerations that make me suspect it is ultimately unsatisfactory as well.

Perhaps the most important is the following. We have noted that a difference in temporal token-reflexivity between perception and thought does not seem to be the right difference to capture the special presence to the perceiving mind of colors, sounds, motion, and so on. But it is unclear in what way *causal* token-reflexivity fares any better here. To see the point, consider that thoughts too can represent their objects as causing them, they just have to do so via their content rather than attitude. It is perfectly possible to think *that* the rectangle is red *and* its being red causes one's thought to that effect. Obviously, however, such content-based causal token-reflexivity does *not* make the thought-about red rectangle present to one 'in the flesh.' So it must be part of the story that there is something special about imposing a causal token-reflexivity requirement *through the attitude* that makes the red rectangle 'present' in the relevant way. It is hard to see why this should be so.

In addition, some work would need to be done to show that the difference between attitude-based and content-based causal token-reflexivity is really first-personally manifest.

Finally, I think the causal token-reflexivity account is on reflection ill motivated. The account is motivated chiefly by the case of veridical hallucinations. Something is *wrong* with these, clearly, but the diagnosis of what is wrong is up for debate. The account under consideration offers us a way to capture theoretically what is wrong with them, by casting them as representationally defective. When the token-reflexivity is built into the content, as in Searle 1983, the relevant experiences turn out to be *non-veridical*; when it is built into the attitude, as in Recanati 2007, they turn out to be incorrect (or inappropriate). But on both accounts, they turn out to be 'representationally defective.' However, an alternative diagnosis would cast them rather as *epistemically* defective. To see why, compare the case of so-called veridical hallucination with the following doxastic case. Suppose a person believes that there is a God, not on the basis of any evidence however, but on the basis of her emotional needs, or because she was brainwashed as a child. And suppose further that, as it happens, there really is a God. This person's belief lacks some of the epistemic qualities highlighted by reliabilists (e.g., it is not *sensitive*: the person's having the belief does not counterfactually depend on the belief's truth). More fundamentally, God's existence is causally unrelated to the person's *having* the belief in God. Nonetheless, intuitively the belief is still true, and therefore correct/appropriate. To cast the belief as incorrect and perhaps false is to mistake the belief's *intentional* character with its *epistemic* character. It should be possible to reach true beliefs by sheer luck! By the same token, I suggest, it should be possible to luck into correct perceptual experiences. And this is just what happens with veridical hallucinations: they are hallucinations because, despite representing the world the way it really is, they are not linked to the world in a way that would bestow on them the relevant positive epistemic property; but they are veridical nonetheless, precisely because they present the world as being a certain way, and the world really is that way. To cast veridical hallucinations as incorrect is thus to confuse their intentional and epistemic profiles.

If the epistemic diagnosis of what is wrong with so-called veridical hallucinations is more accurate than the representational diagnosis, then the causal token-reflexivity account of the intentional character of perceptual experience is unmotivated. As we

have seen, it is also far from clear that the account really delivers a first-personally manifest difference between perception and thought, let alone the kind of difference suited for capturing the distinctive 'in the flesh' character of perception.

With this I close my discussion of attitudinal approaches to the (first-personally manifest) perception/cognition divide. As with the content approaches, I do not pretend to have considered all possible options. For example, I have not considered the notion that perception represents-as-*appearing*-to-obtain whereas thought represents-as-*really*-obtaining (was this Aristotle's view?). I have not considered it essentially because here too the difference does not seem to me to be the right kind to capture the distinctive in-the-flesh character of perception. But the view, as others too, surely deserves treatment in a fully comprehensive examination of possible attitudinal theories of the perception/cognition divide. Still, it is significant that, having considered several prominent or natural content-based and attitudinal approaches, we have found all of them wanting.

7 Conclusion: What Next?

If none of the aforementioned accounts of the first-personally manifest difference between perception and thought is satisfactory, what are our options?

One option is to look for a *non-intentional* difference. We have considered only accounts adverting to intentional character, understood as including both content and attitude. But this leaves open the possibility of an account that adverts neither to content nor to attitude. Perceptual and cognitive states have many other properties, after all. The old idea that perception is experienced as passive or 'receptive' whereas thought is experienced as active or 'spontaneous' may be one example. That particular idea is problematic, as the tendency of thoughts to involve a feeling of agency or authorship may not be a universal characteristic – think of cases where a thought 'pops up' in your mind unbidden ('I haven't paid yet the phone bill!'). But some other altogether non-intentional characteristic(s) may prove essential to perception or thought.

Recall, however, that our concern here is with finding a *first-personally manifest* difference between perception and thought. So what we would need to find are first-personally manifest properties of perception and thought that go beyond its content properties and attitudinal properties. As noted in §2, however, this may run afoul of an extremely modest version of transparency, one which allows not only content but also attitudinal intentional properties to be manifest to introspection. Thus an immediate challenge to the non-intentionalist strategy is to show that perceptual experiences and thoughts *have* altogether non-intentional properties that are first-personally manifest.

A related option is to retain an intentionalist approach to thought but adopt naïve realism about perceptual experience. In this approach, the difference between perceptual experience and thought is that perceptual experience is constitutively related to the perceived object, so that the perceived object is in fact a constituent of the perceptual experience, whereas none of this is true of thought. This thought seems particularly well suited to capturing the distinctive in-the-flesh character of perceptual experience.

The obvious problem for this approach is that envatted brains with mental lives phenomenally indistinguishable from ours will perforce experience the exact same in-the-flesh-ness, despite lacking any kind of constitutive connection to external objects being perceived. What this brings out is that we need to distinguish the *experience* of in-the-flesh encounter with an object, which is a first-personally manifest feature of perceptual states, from *actual* encounter with an object, which goes beyond what is manifest to the subject's first-person perspective. What the naïve realist has to show is that she is well positioned to capture not only actual encounter but also the experience of encounter; for it is only the latter that seems relevant to drawing the first-personally manifest perception/cognition divide.

A third option we have is to *deny the datum*. In a particularly radical form, this would be to deny that there is a substantive distinction between perception and cognition. In a milder form, it might be granted that there is such a distinction, but denied that it has a first-personally manifest dimension. If no account of the first-personally manifest difference between perception and thought seems to work, this may rationalize revisiting the starting assumption that there *is* such a difference.

Perhaps *only* the third-person subpersonal difference in underlying functional architecture is psychologically real.

The immediate challenge for this approach is to explain, or explain away, the fact (as I claim it is) that naïve subjects with no insight into the functional architecture underlying their mental life can still – reliably, effortlessly, immediately, and confidently – determine whether certain of their mental states are perceptual or intellectual.

A fourth option is to go *primitivist* about the first-personally manifest difference between perception and thought. Here the idea is to accept, however reluctantly, that we cannot give an informative account of the perception/cognition divide as it appears to the first-person perspective. Some mental states exhibit a phenomenology of perceptuality, some exhibit a phenomenology of cognitivity, each of us is first-personally acquainted with the difference between the two, and that is all there is to it – there is no way to ‘get underneath’ the difference between perceptuality and cognitivity and anchor it in more fundamental or more specific features of conscious experiences. This sort of no-theory theory is unsatisfactory in an obvious way, as it leaves us with no deeper (personal-level) understanding of a basic bright line in our mental life. Nonetheless we may have to learn to live with this sorry predicament. The pill may be made less bitter if the primitivism at the level of the first-personally manifest is married with a substantive subpersonal story about the mechanisms underlying the instantiation of these *sui generis* features of perceptuality and cognitivity.

A final option, of course, is that there is some content or attitudinal difference between perception and thought, subtler than any considered here, that does anchor the first-personally manifest difference between the two. We just have to employ greater invention to identify it than I have managed here. For my part, I am rooting for this ‘hidden factor’ option – while preparing myself for life as a primitivist.³

References

- Beck, J. 2018. ‘Marking the Perception-Cognition Boundary: The Criterion of Stimulus-Dependence.’ *Australasian Journal of Philosophy*, 96: 319-334.

- Chalmers, D.J. 2004. 'The Representational Character of Experience.' In B. Leiter (ed.), *The Future of Philosophy*. Oxford and New York: Oxford University Press.
- Crane, T. 2009. 'Intentionalism.' In A. Beckermann and B. McLaughlin (eds.), *Oxford Handbook of Philosophy of Mind*. Oxford: Oxford University Press
- Dretske, F.I. 1995. *Naturalizing the Mind*. Cambridge MA: MIT Press.
- Fodor, J.A. 1983. *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Harman, G. 1990. 'The Intrinsic Quality of Experience.' *Philosophical Perspectives* 4: 31-52.
- Heck, R.G. 2000. 'Nonconceptual Content and the Space of Reasons.' *Philosophical Review* 109: 483-523.
- Kriegel, U. 2015a. 'Perception and Imagination: A Sartrean Account.' In S. Miguens, G. Preyer, and C. Bravo Morando (eds.), *Prereflective Consciousness: Sartre and Contemporary Philosophy of Mind*. London: Routledge.
- Kriegel, U. 2015b. 'Experiencing the Present.' *Analysis* 75: 407-413.
- Kriegel, U. Forthcoming. 'Phenomenal Intentionality and the Perception/Cognition Divide.' In A. Sullivan (ed.), *Sensations, Thoughts, and Language: Essays in Honor of Brian Loar*. London and New York: Routledge
- Phillips, B. Forthcoming. 'The Shifting Border Between Perception and Cognition.' *Noûs*.
- Pitt, D. 2004. 'The phenomenology of cognition; Or What is it like to think that P?' *Philosophy and Phenomenological Research* 69: 1-36.
- Recanati, F. 2007. *Perspectival Thought*. Oxford and New York: Oxford University Press.
- Siewert, C. 1998. *The Significance of Consciousness*. Princeton: Princeton University Press.
- Stalnaker, R. 1998. 'What Might Nonconceptual Content Be?' *Philosophical Issues* 9: 339-352.
- Stampe, D. 1987. 'The Authority of Desire.' *Philosophical Review* 96: 335-381.
- Tenenbaum, S. 2009. 'Knowing the Good and Knowing What One is Doing.' *Canadian Journal of Philosophy* 39 (supplement): 91-117.

1 I think we can say this with confidence without a full account of the exact relationship between beliefs, judgments, and thoughts(-that). Some hold that beliefs are never conscious, and thoughts are their conscious counterparts. Others hold that beliefs are dispositional states and thoughts are occurrent manifestations of them. Both these models, and others, suggest that if belief that the

triangle is red represents-as-obtaining the triangle's being red, then thinking that the rectangle is red does so as well. There is an additional question of what *distinguishes* these mental states, but let us set it aside here; the answer may not have to do with attitudinal character.

2 There are of course rather involved philosophical debates on the deep nature of causality (see under: counterfactual dependence, production, manipulation). But those will presumably be more or less extensionally equivalent. So the notion of causality invoked here can be neutral between them.

3 This work was supported by the French National Research Agency's grant ANR-17-EURE-0017, as well as by grant 675415 of the European Union's Horizon 2020 Research and Innovation program. For comments on a previous draft, I am grateful to Géraldine Carranante, Ben Phillips, and Enrico Terrone. I have also benefited from presenting the paper at Columbia University, the University of Luxembourg, the University of Milan, and conferences on phenomenal intentionality (Paris, March 2017), perceptual awareness (Paris, July 2017), and perception and observation (Kirchberg, Austria, August 2017); I am grateful to the audiences there, in particular Davide Bordini, Géraldine Carranante, Marian David, Arnaud Dewalque, Anna Giustina, Gabriel Greenberg, Martin Lin, Tricia Magalotti, Olivier Massin, John Morrison, Takuya Niikawa, Elisa Paganini, Jesse Prinz, Susanna Schellenberg, Enrico Terrone, Alfredo Tomasetta, and Nick Young.