

# *The Shifting Border Between Perception and Cognition*

BEN PHILLIPS

University of the Pacific

## **Abstract**

The distinction between perception and cognition has always had a firm footing in both cognitive science and folk psychology. However, there is little agreement as to how the distinction should be drawn. In fact, a number of theorists have recently argued that, given the ubiquity of top-down influences (at all levels of the processing hierarchy), we should jettison the distinction altogether. I reject this approach, and defend a pluralist account of the distinction. At the heart of my account is the claim that each legitimate way of marking a border between perception and cognition deploys a notion I call ‘stimulus-control.’ Thus, rather than being a grab bag of unrelated kinds, the various categories of the perceptual are unified into a superordinate natural kind (*mutatis mutandis* for the complimentary categories of the cognitive).

## **1. Introduction**

Is there a viable distinction to be drawn between perception and cognition? There certainly *seems* to be a difference in kind between hearing a balloon pop and thinking about the square root of  $-1$ . But common sense is not the only area in which the distinction is gainfully employed. As Firestone and Scholl (2016, 4) observe, the distinction is “woven so deeply into cognitive science as to structure introductory courses and textbooks, differentiate scholarly journals, and organize academic departments.”

Contemporary philosophy of mind is certainly brimming with debates that presuppose a perception/cognition border. Examples include the cluster of debates surrounding the nature of perceptual content; the debate as to whether cognitive states can influence perceptual ones; and traditional debates about the epistemological role of perception. How one goes about sorting the perceptual from the cognitive will clearly have a large role to play in determining one’s stances on these debates.

Unfortunately, despite the high regard in which the perception/cognition distinction is held, there is nothing close to a consensus as to how it should be drawn. And the more neuroscience reveals about the cascade of neural processes that begin once stimuli hit our sensory transducers, the further we seem to be moving away from a consensus. Of course, friends of the distinction will agree about paradigmatic cases: e.g. *hearing* a balloon pop versus *thinking* about imaginary numbers. The problem is that there is a slew of cases in between that are neither obviously perceptual nor obviously cognitive. Such cases include object recognition; certain forms of multi-sensory integration; and some varieties of mental imagery. In tackling problematic cases like these, there are three options: monism, pluralism, and eliminativism.

According to the monist, there is a unique border between perception and cognition. Once we carefully distill the explanatory projects that make this border indispensable, some of the problematic cases listed above may turn out to be determinately perceptual, while others may turn out to be determinately cognitive. As for any residual cases, it may be that they are neither determinately perceptual nor determinately cognitive: theoretical notions are not immune to vagueness.<sup>1</sup>

According to the pluralist, there are multiple ways to mark a border between perception and cognition, each one being tethered to a distinct explanatory goal. For instance, perhaps object recognition qualifies as a perceptual capacity relative to explanatory project *A*, but a cognitive one relative to explanatory project *B*.

The third option—one that has been gaining steam in the recent literature—is eliminativism. According to the eliminativist, we should jettison the distinction between perception and cognition altogether. For instance, in discussing the apparent ubiquity of top-down processes, Gary Lupyán (2015*b*, 585) endorses “a collapse between perception and cognition,” such that “the very question of whether cognition affects perception becomes ill-posed.” Similarly, Andy Clark (2013, 10) sees his predictive processing framework as making “the lines between perception and cognition fuzzy, perhaps even vanishing.” And in discussing problematic cases of the sort mentioned above, Nicholas Shea (2015, 87) suggests that the perception/cognition distinction may be a relic of folk psychology: one that fails to do any “deep explanatory work.”

In what follows, I defend pluralism. At the heart of my account is the claim that each legitimate way of marking a border between perception and cognition invokes a notion I call ‘stimulus-control.’ Thus, rather than being a grab bag of unrelated kinds, the various categories of the perceptual are unified into a superordinate natural kind (*mutatis mutandis* for the complimentary categories of the cognitive)—a point that will be crucial in staving off the threat from eliminativism. I conclude by drawing consequences for those debates in philosophy of mind that presuppose a border between perception and cognition.

## 2. Searching for a Border

Before setting off, it will be helpful to get clear on what exactly we are searching for, and how to tell if we have found it. There are, in fact, two searches one can embark on here. One search involves determining whether notions of the perceptual and the cognitive have explanatory significance in cognitive science. The other search involves taking those natural kinds, if any, uncovered in the first search, and determining how they map onto folk notions.

### 2.1. *Carving the Mind at Its Joints*

What it is for a class to constitute a natural kind is a difficult question. For our purposes, the most important requirement is that its members must share certain distinctive properties the appeal to which gives us scientific explanations and inductive inferences that we wouldn’t otherwise have at our disposal. For instance, the disjunctive class of those beliefs that are possessed by either cat-owners or

film-buffs does not constitute a natural kind because invoking it does not arm cognitive scientists with any novel explanations or inductive inferences.

For those embarking on the second type of search, there is another thorny question looming in the background: namely, what constitutes a suitable mapping between folk notions of perception and cognition, and those relevant kinds that earn their keep in cognitive science?

One key desideratum is to avoid a perception/cognition distinction that is purely stipulative.<sup>2</sup> Take the explanatorily fruitful distinction between bottom-up and top-down processes. Construing this as the perception/cognition distinction would do too much violence to the folk conception. There is certainly nothing in the ordinary understanding of perception that rules out top-down influences of the intra-perceptual variety. Beyond this, views on just how far a folk notion can diverge from the natural kind in question, while still referring to that kind, are highly contentious.<sup>3</sup>

In addressing the analogous question about the notions of representation deployed in cognitive science, William Ramsey enjoins us to

... carefully examine the different notions of representation that appear in cognitive theories, get as clear as possible about just what serving as a representation in this way amounts to, and then simply ask ourselves—is this thing really functioning in a way that is recognizably representational in nature? In other words, instead of trying to compare representational posits against some sort of contrived definition, we can instead compare it directly to whatever complex concept(s) we possess to see what sort of categorization judgment is produced ... To some degree, this means our analysis will depend on a judgment call. If this is less tidy than we would like, so be it. (2007, 10)

In carrying out our search for the border (or borders) that the folk are referring to, Ramsey's basic methodology is the one I propose to use. Having said that, my focus will largely be on the first kind of search: the one that seeks to uncover those notions of perception and cognition that earn their explanatory keep in cognitive science.

Having carried out both kinds of search, my conclusion will be that we should adopt pluralism on both fronts. I will do so by arguing that (i) more than one perception/cognition border earns its explanatory keep in cognitive science, and (ii) the relevant border-markers map relatively neatly onto folk conceptions of the perceptual and the cognitive.

## 2.2. *Narrow Versus Broad Notions of the Perceptual and the Cognitive*

The final issue that needs to be addressed before we set off concerns an ambiguity in usages of the terms 'perceptual' and 'cognitive.' Theorists sometimes invoke *three* mutually exclusive categories of the mental: the perceptual, the quasi-perceptual (e.g. imaginative states), and the cognitive. If we go with this taxonomy then being a non-perceptual state does not entail being a cognitive one. Following Beck (forthcoming), we can call the class of states that includes both cognitive and quasi-perceptual states, *broad cognition*—the complimentary class being *narrow*

*perception*. Similarly, we can call the class of states that includes both perceptual and quasi-perceptual states, *broad perception*—the complimentary class being *narrow cognition*.

Even though they go by different names, notions of the narrowly perceptual and the quasi-perceptual have a home in folk psychology. The following are examples: “He claims he *saw* a UFO, but I suspect he *imagined* it”; “I didn’t *see* their angry faces, but I can *visualize* them”; “I can’t *hear* the kookaburra, but I remember what they sound like.” Implicit in attributions of this sort is a distinction between two kinds of mental states: those that represent the environment via sensory means; and those that merely *resemble* such states.

Notions of the narrowly perceptual and the quasi-perceptual also figure centrally in various domains of cognitive science. For instance, Palmer characterizes the notion of perception that is of interest to vision scientists as follows:

Perception is not merely about an observer’s subjective visual experiences, because we would not say that even highly detailed hallucinations or visual images would count as visual perception . . . it is part of visual perception only when it signifies something about the nature of external reality. (1999, 6)

Similarly, in homing in on the notion of perception that earns its keep in perceptual psychology, Burge excludes sensory imagination from the realm of the perceptual, characterizing it as follows:

Sensory imagination is not committal on whether at least some aspects of its representational contents have *representata*, past, present, or future. (2010, 378)

For instance, suppose I visualize a fictional landscape before I start my painting. My visualization does not have the function of representing objects in the scene before me. In contrast, perceptual states “necessarily represent what is perceived—the referents, or purported referents, of perception—as being a certain way” (Burge 2010, 380). They necessarily “function (representationally) to single out concrete particulars” (2010, 380).

I elaborate on the explanatory roles distinctive of narrow perception in sections 3 and 4. For the moment, the important point is that, *prima facie*, there are good reasons for construing narrow perception as a natural kind in its own right.<sup>4</sup>

In addition to the fact that narrow perception plays a distinctive explanatory role, it is worth considering another key rationale for focusing on the distinction between narrow perception and broad cognition. Take the debate surrounding the reach of cognitive penetration. Parties to this debate agree that imaginative states are cognitively penetrable. For example, my belief that it is about to rain might cause me to visualize running for shelter. It follows that the debate surrounding whether ‘perception’ is penetrable must be about top-down influences on *narrow* perception. For instance, in setting up the terms of the debate, Zeimbekis and Raftopoulos (2015) distinguish between narrow perception and “cognitive states like propositional attitudes, concepts, long-term memories, imagery, or cognitively driven attention” (2015, 27). Similarly, in defending their thesis that there is no compelling evidence for the existence of cognitive penetration, Firestone and

Scholl (2016, 65) rule out top-down effects on mental imagery, explicitly drawing a distinction between visual imagery and visual *perception*.<sup>5</sup>

The debate concerning whether ‘perceptual’ content is singular (e.g. <That is *F*>) or purely general (e.g. <There is an *F*>) is also about *narrow* perception.<sup>6</sup> Parties to this debate presumably agree that when I visualize a fictional landscape, my visualization does not aim to single out environmental particulars. Thus, in disagreeing about ‘the contents of perception,’ they are presumably disagreeing about the contents of *narrow* perception.

Getting clear on what marks the border between narrow perception and broad cognition thus promises to shed light on the debates outlined above. But it also provides further motivation for the view that narrow perception is an important kind: one whose explanatory roles diverge from broad perception.

In what follows, I use the term ‘perception’ to refer to narrow perception, and ‘cognition’ to refer to broad cognition. As will become clear, my core thesis is that any marker of the perceptual must deploy a notion of stimulus-control.

### 3. Stimulus-Control

Perhaps the most obvious way to go about marking a border between perception and cognition is to deploy the notion of stimulus-dependence. When I close my eyes and block my ears, I fail to see or hear the popping of the balloon. But I can certainly *think* that balloons make a popping sound without registering stimuli with my eyes or ears. Perceptual states seem to be stimulus-dependent in a way that cognitive states are not. Let’s define stimulus-dependence as follows:

#### *Stimulus-Dependence*

A state or process is stimulus-dependent just in case it causally depends on the transduction of a proximal stimulus.<sup>7</sup>

#### 3.1. A Functional Approach

An obvious problem with the bare claim that perceptual states are stimulus-dependent is that thoughts are often caused by the transduction of a proximal stimulus. For example, suppose my thought that balloons are fun was caused by the proximal array of light reflected from the balloon I am holding. If we mark a perception/cognition border solely in terms of stimulus-dependence, this thought will qualify as perceptual, when clearly it is not.

An obvious solution is to characterize perceptual states in functional terms: in contrast to cognitive states, perceptual states have the *function* of representing the environment in a stimulus-dependent fashion (see Beck (forthcoming)). My thought that balloons are fun may well have been caused by a proximal stimulus, but this fact about its etiology is not constitutive of its function.

There are various ways in which we might elucidate the notion of a function. According to the etiological approach, a trait of an individual has function, *F*, if and only if its performing *F* contributed to the survival and reproductive success of the individual’s ancestors.<sup>8</sup> The main competitor to the etiological approach is

the propensity theory, according to which the function of a trait is that which *will* contribute to the survival and reproductive success of the individual in question.<sup>9</sup>

I am not wedded to a specific theory of functions. Whichever approach we adopt, the fact remains that narrowly perceptual states perform functions that broadly cognitive ones do not. I elaborate on some of these functions below. In brief, though, we have already seen that one of the most obvious differences between narrow perceptions and broad ones is that the latter need not aim to represent the actual environment.

### 3.2. *Being Causally Controlled by the Stimulus*

The functional characterization given above retains the intuitive idea that stimulus-dependence has a central role to play in distinguishing the perceptual from the cognitive. As the following remarks attest to, it is also a characterization that has currency among perceptual psychologists:

[T]he word ‘perception’ will be used throughout this article to refer to the extraction and/or representation of perceptual information from a stimulus (Kanwisher 2001, 90).

In the context of this book, visual perception will be defined as the process of acquiring knowledge about environmental objects and events by extracting information from the light they emit or reflect. (Palmer 1999, 5)

The visual system isn’t interested in reporting on the nature of the retinal image. Its task is to use retinal images to deliver a representation of what is *out there in the world*. (Frisby & Stone 2010, 16)

These characterizations are all clustered around the core thesis that the function of a perceptual state is to represent environment entities in a way that is causally mediated by the stimuli they produce.

Nonetheless, we are still saddled with a problem. Suppose I saw a bonobo in the Congo Basin yesterday. Today, I have the thought, *I SAW A BONOBO YESTERDAY*. This thought arguably functions so as to represent the bonobo in a way that causally depends on the proximal stimulus she produced. And yet, we would not want to say that my thought is a perceptual state.

In distinguishing between thoughts about past perceptions, and the perceptions themselves, the key notion is *causal control*: perceptual processes function so as to be causally controlled by sequences of proximal stimuli, whereas, cognitive processes do not.<sup>10</sup>

The notion of causal control figures centrally in theories of genetic coding. For instance, genes are often construed as causally controlling the sequencing of amino acids (Godfrey-Smith 2000, 5; Stegman 2014). In elucidating the notion of causal control, Stegman appeals to two types of causal structures. An “internal ordering” is a sequence of effects,  $\langle F, G, H \rangle$ , such that each member was brought about by its predecessor ( $F$  caused  $G$ , and  $G$  caused  $H$ ). On the other hand, an “external ordering” is a sequence of effects,  $\langle X, Y, Z \rangle$ , such that each member was caused by events that are not part of the sequence ( $A$  caused  $X$ ,  $B$  caused  $Y$ , and  $C$  caused  $Z$ ).

As Stegman (2014, 452–5) points out, any process that constitutes an external ordering qualifies as being causally controlled by the external process in question. For instance, the rolling film of a projector causally controls the sequence of images on the screen; the movement of the pianist’s fingers causally controls the sequence of tones emitted by the piano’s strings; and the order in which amino acids are produced is under the causal control of certain genes.

Applying this to perception, our new characterization becomes this: a process is perceptual just in case it has the function of producing representations of environment entities by being causally controlled by those proximal stimuli that these entities produce. Call this thesis, *Stimulus-Control<sub>P/C</sub>*.<sup>11</sup> For instance, my current thought that I saw a bonobo yesterday is not under the causal control of the ongoing stream of proximal stimuli. I can maintain this thought during an episode in which I close my eyes for 30 seconds, open them, and then swivel 180° to look out the window. The stream of proximal stimuli that span this episode need not exert any causal control over my thought. On the other hand, if this stream of proximal stimuli does not causally control the perceptual processes that ensue, a malfunction will occur: one that manifests itself in a visual illusion, and a failed action (e.g. bumping into the window).

### 3.3. *Stimulus-Control and Top-Down Processes*

In clarifying the thesis that perceptual processes are stimulus-controlled, it is important to realize that it does not rule out top-down influences. The characterization given above merely says that perceptual processes function so as to be causally controlled by proximal stimuli: it does not say that proximal stimuli are always the *only* causal factors in play.

In fact, even if a paradigmatically perceptual process is itself top-down, this does not entail that it is not stimulus-controlled. Feedback mechanisms in early visual areas (e.g. from V2 to V1) function so as to be causally controlled by proximal stimuli. Without stimuli, there is nothing for an intra-perceptual feedback mechanism to operate on (see Kafaligonul *et al.* 2015).

### 3.4. *The Explanatory Significance of Stimulus-Control*

The way in which we just marked the perceptual off from the cognitive is clearly wedded to what perceptual psychologists call *the underdetermination problem*.<sup>12</sup> The problem is that for any given proximal stimulus, there is a multitude of environmental scenes that could have caused it. For example, a convex object illuminated from above, and a concave object illuminated from below, can produce the same pattern of light on the retina.

In explaining how the visual system solves the underdetermination problem, perceptual psychologists posit what Burge (2010, 345–6) calls *formation principles*. These principles describe law-like regularities in which proximal stimuli are transformed into representations of their distal causes. Given that these transformational regularities mirror law-like regularities in the distal environment, perceptual systems represent causes of proximal stimuli in a relatively veridical fashion. For example, the visual system tacitly assumes that light tends to come from above—an

assumption that has a large role to play in determining why craters are normally perceived as such, and not as mounds, illuminated from below.

The transformations that resolve instances of the underdetermination problem are clearly under the causal control of proximal stimuli. How could it be otherwise? For instance, how could one's visual system overcome the underdetermination problem and represent a moving object, unless the series of proximal stimuli caused by this object were causally controlling the ensuing transformations?

Explaining how perceptual systems solve instances of the underdetermination problem is at the heart of perceptual psychology (see Burge 2010, 345). Given that the processes we have marked off as perceptual are just those processes that function so as to resolve underdetermination problems, it follows that the perception/cognition border we have uncovered has explanatory significance in perceptual psychology.

Now that I have clarified the notion of stimulus-control, the plan for the rest of the paper is as follows. First, I introduce another traditional way of marking a perception/cognition border—one that invokes the notion of modality-specificity. I argue that one fruitful way of cashing out the notion of modality-specificity makes it a species of stimulus-control. I then address various other attempts to mark a border between perception and cognition. In each case, I argue that unless we invoke stimulus-control, the border that emerges is not one that we should recognize as demarcating the perceptual from the cognitive.

#### 4. Modality-Specificity as Stimuli-Specific-Control

Suppose you're having an audiovisual experience of a wahooing baboon. It is plausible that distinct aspects of your experience are tied to distinct sensory modalities: the representations of the baboon's color, shape, and motion are tied to your visual system; while the representations of the pitch and volume of its wahooing are tied to your auditory system. In other words, each aspect of your experience is *modality-specific*.

According to one traditional way of marking a border between perception and cognition, perceptual processes are modality-specific, whereas cognitive states are amodal (see McFarland and Cacace 1999; Bayne 2014). For instance, in contrast to my auditory perception of the baboon's vocalization, my thought that the square root of  $-1$  is an imaginary number can be tokened in the absence of any given stimulus-type. The notion of being tied to a specific sensory modality is intuitive enough, but what exactly is a sensory modality, and what exactly does being "tied to" one entail?

##### 4.1. Determining the Upper-Border of a Sensory Modality

In determining what counts as a sensory modality, we need to distinguish between two questions: (1) What distinguishes the senses? (2) What distinguishes a given sensory modality from cognition? The first question is about horizontal borders. The second one is about vertical borders. Importantly, standard answers to (1) often leave (2) open.

Take the Aristotelian thesis that for each sensory modality, there is at least one property that it uniquely represents (*De Anima* 418a11-16). For instance, vision is the only modality that represents colors; audition, the only modality that represents sounds; and so on. Even if this way of marking the borders between the senses is viable, it does not settle the question of where to mark those borders that separate each sensory modality from cognition—cognitive states can certainly represent colors and sounds.

Similar remarks apply to the claim that each modality produces states with a distinctive phenomenology. That may well constitute a plausible answer to the question about horizontal borders, but it does not settle the issue of where each sensory modality ends, and cognition begins (see Macpherson 2011).

No doubt, one could supplement standard answers to (1) in a way that provides answers to (2). For example, one might mark the upper-borders of the sensory modalities by noting that modality-specific states and cognitive states possess contents (or phenomenology) of a fundamentally different *kind*. I consider both content-based and phenomenological approaches to the perception/cognition distinction in section 7, where I argue that an appeal to stimulus-control is still mandatory.

This leaves us with two other standard ways of elucidating *modality-specificity*: the appeal to proprietary sense-organs; and the appeal to proprietary stimuli. In pursuing the sense-organ approach, we would need to answer questions about what makes something an eye; what makes something an ear; and so on. And it is very hard to think of plausible answers that do not invoke the types of stimuli that are unique to each organ: eyes process light; ears process pressure waves; etc. The sense-organ account appears to collapse into the proximal stimulus account.

#### 4.2. *Stimulus-Specific-Control*

According to the proximal stimulus account, the way to individuate sensory modalities is to distinguish between those processes that depend on a specific type of stimulus versus those that do not. For example, your auditory experience of the baboon's wahooing causally depends on the pressure waves being transduced by receptors in your inner ear, while your visual experience of the baboon's mouth causally depends on the array of light being transduced at your retinas. Modifying *Stimulus-Control*<sub>P/C</sub>, this gives us the following:

##### *Stimulus-Specific-Control*<sub>P/C</sub>

A process or state is perceptual just in case it functions so as to represent environmental entities in a way that is causally controlled by a specific type of stimulus.<sup>13</sup>

Below, I explain how *Stimulus-Control*<sub>P/C</sub> and *Stimulus-Specific-Control*<sub>P/C</sub> give us different verdicts on certain cases. Before I do that, two key aspects of stimulus-specific-control are worth noting. First, it is clearly a sub-type of stimulus-control—a point I return to in dealing with the threat from eliminativism. Second, it is an explanatorily fruitful notion. Those states that exhibit stimulus-specific-control

just are those states that are constitutively involved in resolving stimulus-specific instances of the underdetermination problem.

Moreover, if we know that *A* is undergoing a mental state that is causally controlled by the transduction of light, but not sound, while *B* is undergoing a state that is causally controlled by the transduction of sound, but not light, that enables us to generate different explanations of their respective behaviors. On the other hand, if all we know is that *A* and *B* are undergoing stimulus-controlled states, we won't have the same predictive and explanatory purchase on their behaviors.

#### *4.3. Cross-Modal Integration and Stimulus-Specific-Control*

One threat to the explanatory power of stimulus-specific-control comes from a consideration of cross-modal influences. When you see someone's lips move while hearing them talk, you don't just have a visual experience and an auditory one at the same time. Rather, you undergo a multisensory experience in which you perceive the sounds as emanating from the lips. This is corroborated by studies in which altering the motion of the lips, whilst keeping the sound constant, alters the subject's auditory experience. For instance, in the McGurk effect, a "ba" sound is played while the subject sees the lips of the speaker making a "ga" movement, the result being an experience of a "da" sound (McGurk and MacDonald 1976). If cross-modal influences of this sort are ubiquitous—and the evidence suggests that they are (Nudds 2015; Shams and Kim 2010)—doesn't this undermine the claim that paradigmatically perceptual states, such as seeing and hearing someone speak, are causally controlled in a stimulus-specific manner?

The existence of cross-modal influences does not undermine the explanatory status of stimulus-specific-control per se. Just as the intervention of top-down influences does not preclude a process from qualifying as stimulus-controlled, the intervention of lateral influences (of the cross-modal variety) does not preclude a process from being controlled in a stimulus-specific manner. Reconsider the McGurk effect. One experiences the "da" sound, in part, because of the proximal array of light transduced at one's retinas. However, the auditory experience of the "da" sound could have occurred in the absence of retinal stimulation without that experience thereby constituting a malfunction. The experience of the "da" sound would only qualify as a malfunction if it were tokened in the absence of a causally controlling auditory stimulus.

A more pressing concern with the explanatory status of stimulus-specific-control involves states that are paradigmatically perceptual, but have the function of being causally controlled by multiple stimulus-types. Take flavor perception. Normally, when I experience something's flavor, there are several contributing factors, such as its taste; its texture; and its smell, which is registered in the nose via retronasal olfaction (Auvray & Spence 2008). What is so interesting about flavor perceptions is that their occurrence seems to require a confluence of stimulus-types. The easiest way to demonstrate this to yourself is to bite into a block of Camembert cheese, before taking another bite with your nose blocked. By subtracting the influence of the cheese's smell, the original percept is supplanted. If flavor perception really does

work in this way, doesn't that spell trouble for the thesis that a perceptual state is causally controlled by a specific type of stimulus?

Fortunately, there is a fix available. Let's say that a process or state is causally controlled in a *stimuli-specific* manner just in case it is causally controlled by a specific *combination* of stimulus-types (those processes that are causally controlled by a single type of stimulus being the degenerate cases). Thus, according to our amended account, perceptual states are those that function so as to be causally controlled in a stimuli-specific manner: call this thesis, *Stimuli-Specific-Control<sub>P/C</sub>*.

For example, my experience of the Camembert's flavor is perceptual according to our present approach, because it functions so as to be causally controlled by gustatory and olfactory stimuli, but not other types, such as pressure waves or light.

I have by no means provided definitive arguments for the thesis that, in addition to those states that function so as to be causally controlled in a stimulus-specific manner, there are also states that function so as to be causally controlled in a *stimuli-specific* manner. However, in order to motivate pluralism, doing so is not necessary: either way, there is an explanatorily significant border to be marked between those states the function of which is to be causally controlled by either some unique stimulus-type or combination of stimulus-types, and those states that do not possess this function. And as I will now argue, even though there are points of convergence, there are points at which this border diverges from the one marked according to *Stimulus-Control<sub>P/C</sub>*.

## 5. *Stimulus-Control<sub>P/C</sub> Versus Stimuli-Specific-Control<sub>P/C</sub>*

Many of the capacities that manifest themselves early on in the processing hierarchy are rendered perceptual regardless of whether we go with *Stimulus-Control<sub>P/C</sub>* or *Stimuli-Specific-Control<sub>P/C</sub>*: edge-detection, amodal completion, and color-perception are all examples. Take visual amodal completion: the capacity to see objects as having occluded parts. Unless there has been a malfunction, proximal stimuli produced by visible object-parts causally control visual completion processes. Moreover, they do so in a stimulus-specific manner: visual completion is causally controlled by retinal images, not by auditory stimuli, or any other type of stimulus. Visual amodal completion thus qualifies as perceptual regardless of whether we go with *Stimulus-Control<sub>P/C</sub>* or *Stimuli-Specific-Control<sub>P/C</sub>*.

### 5.1. *Perceptual Memory and Anticipation*

What do *Stimulus-Control<sub>P/C</sub>* and *Stimuli-Specific-Control<sub>P/C</sub>* say about those states that (i) are necessary for solving instances of the underdetermination problem, but (ii) are not themselves representations of the current environment?<sup>14</sup> Take visual short-term memory (VSTM)—the capacity to briefly retain a very limited amount of visual information (usually, 3–4 items).<sup>15</sup> VSTM is necessary for paradigmatically perceptual processes to occur. For example, perceiving an object in motion requires the coordination of VSTM and perceptual anticipation (Burge 2010, 447). In fact, when visually tracking a stationary object, each saccade generates a fresh instance of the underdetermination problem: namely, *was the current retinal image caused by*

*the same object that caused the preceding image?* VSTM clearly has a constitutive role to play in solving this type of underdetermination problem (see Hollingworth *et al.* 2008).

As Burge (2010, 378) points out, visual short-term memories and perceptual anticipations are not controlled by proximal stimuli in the same “immediate” way that perceptions of the *current* environment are. Rather, they are formed *from* perceptions of the current environment (Burge 2010, 378). For this reason, Burge denies that states of VSTM and anticipation are perceptual: for him, they are of a piece with sensory imagination and other quasi-perceptual states (2010, 378).

Burge is right that the notion of perception he has homed in on constitutes an important psychological kind. Call the states he is concerned with, *very narrow perceptions*. Very narrow perceptions function so as to represent the *current* environment in a way that is causally controlled by proximal stimuli.

On the other hand, we shouldn’t elide the distinction between quasi-perceptual states, such as voluntary visualizations, and those states that have an eliminable role to play in solving instances of the underdetermination problem, such as visual short-term memories. Even though the latter do not function so as to represent the *current* environment, they function so as to represent the environment in the immediate past or future. Moreover, they do so in a way that is causally controlled by the sequence of proximal stimuli.

For example, visual short-term memories will not aid me in tracking my coffee mug across saccades unless these memories are causally controlled by the sequence of retinal images that these saccades generate. Any short-term visual memory that is not causally controlled in this manner will count as a malfunctioning state: one that will cause me to perform erroneous actions, such as spilling my coffee. Similar remarks apply to perceptual anticipation.

Given that visual short-term memories and anticipations function so as to represent the environment in a way that is causally controlled by proximal stimuli, they qualify as perceptual if we deploy either *Stimulus-Control*<sub>P/C</sub> or *Stimuli-Specific-Control*<sub>P/C</sub>. In contrast, voluntary visualizations and other quasi-perceptual states do not count as (narrowly) perceptual, regardless of which border-marker we use.

In seeking to demarcate the perceptual from the cognitive, I see no reason to privilege the very narrowly perceptual over the narrowly perceptual: both constitute important kinds. Pluralism is thus already on the cards, even before we have located points of divergence between *Stimulus-Control*<sub>P/C</sub> and *Stimuli-Specific-Control*<sub>P/C</sub>. Moreover, the arguments for pluralism—presented below—all focus on points of divergence regarding very narrow perception. Thus, even if one were to build a case for privileging the very narrowly perceptual, pluralism would still be on the cards.

### 5.2. *Hallucinations and Visualizations*

Hallucinations are states that have not fulfilled their function of representing the environment in a way that is causally controlled by proximal stimuli. They therefore count as perceptual according to *Stimulus-Control*<sub>P/C</sub>. For instance, suppose a mad scientist directly stimulates neurons in my visual cortex and generates a visual hallucination as of a *Tyrannosaurus rex*. Even though my hallucination is not under the

causal control of proximal stimuli, it still has the function of representing the environment in a stimulus-controlled manner. Moreover, given that my hallucination functions so as to be causally controlled by *visual* stimuli, it qualifies as perceptual according to both *Stimuli-Specific-Control*<sub>P/C</sub> and *Stimulus-Control*<sub>P/C</sub>.

In contrast to hallucinations, voluntary images do not even function so as to represent the environment in a stimulus-controlled manner. They are not states that have a constitutive role to play in solving instances of the underdetermination problem. Voluntary imagery therefore qualifies as cognitive regardless of whether we deploy *Stimulus-Control*<sub>P/C</sub> or *Stimuli-Specific-Control*<sub>P/C</sub>.

### 5.3. Multisensory Object Files

So far, we have focused on points of convergence between *Stimulus-Control*<sub>P/C</sub> and *Stimuli-Specific-Control*<sub>P/C</sub>. A compelling point of divergence concerns object perception. It is widely held that there is a relatively early stage of object perception in which three-dimensional bodies are individuated and tracked in a stimulus-controlled fashion (see Pylyshyn 2003). One popular way of construing this stage is to invoke *object files*: representations that maintain referential contact with objects through changes in color, shape, location, and so on (Kahneman, Triesman, & Gibbs 1992). As more information about the object comes in from various sensory channels, the file is updated.

Evidence for *visual* object files comes from a phenomenon called the *object specific preview benefit* (see Palmer 1999). Two objects appear on a screen, with letters briefly displayed on them. Each object then moves to a new location. In *same object trials*, the same letter appears on an object before and after it moves. In *different object trials* letters are swapped between objects. Finally, in the third trial, completely novel letters appear. It turns out that subjects are quickest at naming the letter in the same object trial, the explanation being that visual information about an object is stored in a single file.

Now, this provides evidence for *visual* object files, but as Nudds (2015, 184) points out, there are multi-modal incarnations of the object specific preview benefit. In the study carried out by Jordan, Clark, and Mitroff (2010), subjects were shown two objects with pictures on them. The objects then moved to a new position, where a sound was played. In the same object trial, the sound matched the picture that was previously displayed on the given object (e.g. a dog barking, if the initial picture was a dog). Again, an object specific preview benefit was found, strongly suggesting that object files group information from multiple sensory channels into what Nudds (2015, 184) calls “amodal object files.”

Are these amodal object files perceptual or cognitive? If we go with *Stimulus-Control*<sub>P/C</sub> then they are perceptual through and through: tracking an object clearly requires one to token a sequence of states that represent that object in a way that is causally controlled by the sequence of proximal stimuli it produces.

If, however, we go with *Stimuli-Specific-Control*<sub>P/C</sub>, amodal object files appear to be hybrid states. In the experiment described above, both auditory and visual information enter into the object’s file, but this does not undermine the fact that, in order to enter the file, the visual information relies on light transduction, while

the auditory information relies on the transduction of pressure waves. Amodal object files thus appear to represent *object-features* by being causally controlled in a stimuli-specific manner. On the other hand, a file maintains referential contact with its object by being under the causal control of proximal stimuli, but it does so in a stimuli-non-specific manner. For instance, I can update information about the helicopter I am tracking by seeing it; then by hearing it; then by seeing it again. Amodal object files thus appear to be stimuli-specific feature-representations bound to a stimuli-non-specific (or ‘amodal’) singular element. The upshot is that if we go with *Stimuli-Specific-Control<sub>P/C</sub>*, amodal object files have both cognitive and perceptual elements; whereas, if we go with *Stimulus-Control<sub>P/C</sub>*, they are perceptual through and through.

#### 5.4. *Perceptual Demonstrative Thoughts*

Now consider perceptual demonstrative thoughts. Suppose I look at the cat meowing in the tree, consciously inferring that it is stuck. I token the thought, *THAT CAT NEEDS TO BE RESCUED*. My thought’s demonstrative element is causally controlled by the proximal stimuli I receive: if I close my eyes, and block my ears, I can no longer sustain it. In contrast, my thought’s attributive element is clearly not under the causal control of proximal stimuli. It was consciously inferred from background beliefs, and could have occurred as part of a non-demonstrative thought, such as *SOMETIMES CATS NEED TO BE RESCUED*. Demonstrative thoughts with attributive elements of this sort are thus hybrid perceptual-cognitive states if we go with *Stimulus-Control<sub>P/C</sub>*. If, however, we go with *Stimuli-Specific-Control<sub>P/C</sub>*, they are cognitive through and through, for as was the case with amodal object files, the demonstrative element is not under the causal control of any specific combination of stimulus-types.<sup>16</sup>

#### 5.5. *Rapid Object Recognition*

*Needing to be rescued* is presumably not a perceivable category. More generally, given that object recognition involves subsuming objects under concepts stored in long-term memory, many theorists are quick to construe it as a cognitive capacity (e.g. see Firestone and Scholl 2016, 41). But certain instances of object recognition also exhibit some of the hallmarks of paradigmatically perceptual capacities. For instance, object recognition can be extremely swift (Keyesers *et al.* 2001; Grill-Spector and Kanwisher 2005; Potter *et al.* 2014); some forms arguably exhibit adaptation effects (Block 2014); and some processes of object recognition are automatic and non-revisable (Bayne 2009; Mandelbaum (forthcoming)).

Consider swiftness. Potter *et al.* (2014) found that subjects can successfully categorize objects that are presented for only 13 ms, followed by a mask. In fact, in their study, Grill-Spector and Kanwisher (2005) found that subjects are sometimes equally quick at categorizing objects and merely detecting their presence.

As Mandelbaum (forthcoming) points out, these findings are important, for it takes at least 50 ms for the shortest reentrant loops to be completed within vision (Potter *et al.* 2014). It follows that there is just not enough time for

top-down processes to influence those instances of object recognition described above. Mandelbaum draws the following conclusion:

[T]he processing appears to be accomplished via a mental process that is super fast acting, whose processing starts automatically once it receives its proper input, and processes ballistically and without the need of awareness . . . (forthcoming, 7)

This fits with the thesis—defended by theorists, such as Bayne (2009)—that certain instances of object recognition are perceptual because they are insensitive to countervailing evidence. I discuss that approach to the perception/cognition distinction in the next section.

For present purposes, the key point is that rapid recognition appears to serve a distinctive function. Potter and colleagues (2014) characterize this function as follows:

A possible role for such rapid visual understanding in normal experience would be to provide nearly instantaneous conceptual activation that enables immediate action when necessary, without waiting to refine understanding by reentrant processing or by the kind of conscious reflection that requires a stable recurrent network. (2014, 278)

Mandelbaum (forthcoming, 10) elaborates on this suggestion, arguing that having a special system for *rapid* recognition makes good evolutionary sense. For instance, when it comes to panther identification, making an all-things-considered judgment about whether *X* is a panther would be time-consuming and dangerous.

One might wonder whether the concepts tokened as a result of rapid recognition can also figure as constituents of paradigmatically cognitive states. For example, I can obviously have a thought about panthers without perceiving one. That may well be true, but it does not follow that the outputs of rapid recognition violate *Stimulus-Control*<sub>P/C</sub>. All that follows is that rapid recognition can produce conceptualizations that are type identical to those that figure in cognition. This is entirely compatible with the thesis that rapid recognition functions so as to produce object representations in a way that is causally controlled by proximal stimuli. If rapid recognition did not fulfill this function, it would not have the payoffs outlined above.

There are thus good reasons for thinking that rapid recognition satisfies *Stimulus-Control*<sub>P/C</sub>. What is the verdict of *Stimuli-Specific-Control*<sub>P/C</sub>? In recognizing a piano *as such*, I might do so by seeing one; but I could also recognize one by feeling its keys or by hearing someone play it. Piano-recognition is not causally controlled by any specific stimulus-type or combination of stimulus-types. Neurophysiological evidence supports this point: several studies have revealed populations of neurons that selectively respond to specific types of objects, regardless of the stimulus-type used to perceive them (Quiroga *et al.* 2009; Man *et al.* 2012; Yildirim and Jacobs 2012). Rapid object recognition is thus a good candidate for a capacity that is perceptual if we go with *Stimulus-Control*<sub>P/C</sub>, but cognitive if we go with *Stimuli-Specific-Control*<sub>P/C</sub>.<sup>17</sup>

Let's take stock: in carrying out our border search, we have uncovered two joint-carving approaches: *Stimulus-Control*<sub>P/C</sub> and *Stimuli-Specific-Control*<sub>P/C</sub>. In surveying just some of the verdicts given by these two approaches, we have found points of agreement, and points of divergence. In what follows, I consider various other attempts to mark a perception/cognition border. In each case, I argue that an appeal to stimulus-control is unavoidable.

### 6. Impenetrability as the Mark of the Perceptual

The thesis that perceptual systems are modular input-analyzers was inspired by the work of Fodor (1983) and Pylyshyn (1984). At present, it is widely held that paradigmatically perceptual processes do not meet Fodor's central criteria for modularity. Take the criterion of encapsulation. System *A* is encapsulated relative to system *B* just in case *A* does not have access to information stored in *B*. It is now widely agreed that, given the ubiquity of cross-modal influences, perceptual systems are not fully encapsulated relative to one another. The thesis that perceptual processes are not encapsulated relative to so-called central cognition is far more contentious.<sup>18</sup>

Let's sidestep this controversy and suppose that there really is a boundary to be drawn around those processes that cannot be penetrated by paradigmatically cognitive states. Would this thereby furnish us with a legitimate way to mark a border between perception and cognition?

In order to address this question, imagine an abnormal subject whose brain has been reengineered in such a fashion that paradigmatically cognitive processes frequently influence paradigmatically perceptual ones. In fact, let's imagine that top-down influences of this sort are rampant. Would this thereby remove any border between the perceptual and the cognitive for this individual's mental states? As I mentioned above, this appears to be the view of theorists, such as Clark (2013) and Lupyan (2015a, 2015b). But I fail to see how the argument even gets off the ground.

If cognitive penetration does occur, there must be significant limitations on its reach. For instance, if an ancestor's desire to avoid danger (e.g. predators) regularly caused her to visually experience an absence of such dangers—regardless of the nature of the stimulus—that individual would not enjoy much survival and reproductive success. For top-down effects to have survival and reproductive value, they clearly need to work in concert with the stimulus.

The upshot is that even if you countenance the types of top-down processes that are often construed as instances of cognitive penetration, you still have to recognize the border between those states that are causally controlled by proximal stimuli, and those that are not: without the former, there would be nothing for top-down processes to operate on. This, in turn, means that you still have to recognize a way of fixing a border between perception and cognition that does not hang on the issue of whether encapsulation obtains, for being influenced by top-down processes does not preclude a state from satisfying *Stimulus-Control*<sub>P/C</sub>.

At the very least then, the absence of encapsulation would not entail eliminativism about the perception/cognition distinction. On other side of the coin,

monism about the perception/cognition distinction would not follow from the *truth* of the encapsulation thesis. Again, to establish that conditional, the proponent of encapsulation would have to undermine the thesis—defended above—that *Stimulus-Control<sub>P/C</sub>* and *Stimuli-Specific-Control<sub>P/C</sub>* are also legitimate border markers.

Most importantly, even if there is a viable distinction to be drawn between those systems that are encapsulated and those that are not, being encapsulated is not sufficient for a process to qualify as perceptual. The rise of the massive modularity thesis is a testament to this fact, with theorists such as Cosmides and Tooby (1992) and Carruthers (2006) providing general arguments for the thesis that processes normally attributed to “central cognition” are encapsulated to a significant degree. Examples include processes of belief- and desire-formation; decision-making processes; and language production.

In dealing with this problem, an obvious solution is to say that an encapsulated process only qualifies as perceptual if it also satisfies *Stimulus-Control<sub>P/C</sub>*. On this way of marking a perception/cognition border, we can think of those processes on the perceptual side of the divide as exhibiting a pure form of stimulus-control: one in which certain types of top-down effects are completely absent.

It is beyond the scope of this paper to assess whether the class of encapsulated-cum-stimulus-controlled processes is co-extensive with either the class of stimulus-controlled processes or the class of processes that are controlled in a stimuli-specific manner—it is controversial whether top-down effects of the relevant sort obtain at all, let alone which specific varieties do. Suffice it to say that there is no a priori reason to expect a complete convergence with either *Stimulus-Control<sub>P/C</sub>* or *Stimuli-Specific-Control<sub>P/C</sub>*, and for obvious reasons, there cannot be a complete convergence with both.

## 7. Phenomenology, Content, and Format

In addition to the candidates we have examined so far, there are three other features that are traditionally appealed to in marking a border between perception and cognition: phenomenology, content, and format.

### 7.1. Phenomenology

The thesis that perceptual states have a proprietary sort of phenomenology is a seductive one. Just consider the stark difference in what it’s like to bite into a lemon versus what it’s like to think about prime numbers. In fact, rather than observing a difference in phenomenal character, many theorists regard states of the latter variety as being devoid of phenomenal character altogether (see Bayne and Montague 2011). Either way, it is widely agreed that there is a border to be marked here. But is it recognizable as a perception/cognition border? One reason for thinking that the appeal to differences in phenomenology will not suffice is that some mental states function in a way that makes them (broadly) cognitive, even though they seem to possess the same type of phenomenology as paradigmatically perceptual experiences. Let me offer two such examples: voluntary visualizations and episodic memories.

If we use introspection as a reliable source of evidence concerning differences in phenomenology, visualizations are plausibly classified as having the same type of phenomenology as paradigmatically perceptual experiences: a claim that is bolstered by the fact that subjects often fail to distinguish between paradigmatic perceptions and visual imagery. For example, in Perky's (1910) famous study, subjects were asked to fixate on the center of a white screen and imagine various objects as being located there. Unbeknownst to them, various faint images were projected onto the point of fixation, just above the threshold of visibility. Even though subjects did not notice the projected images as such, they had a striking influence on their introspective reports. One subject reported trying to visualize a horizontal banana, but was surprised to find that she ended up visualizing a vertical one instead. In general, subjects failed to distinguish between those aspects of their overall experience that were contributed by imagination, and those that were contributed by veridical perception.

Studies of this sort provide evidence that visualizations and paradigmatically visual perceptions share the same type of phenomenology. But should visualizations therefore be construed as perceptual states? Reconsider my visualization of the fictional landscape. This state does not have the function of representing the actual world. As was argued above, this thereby excludes it from the realm of the (narrowly) perceptual. We therefore cannot fix a border between perception and cognition by invoking phenomenal features alone.

In salvaging a border-marking role for phenomenology, a natural move to make at this point would be to couple the appeal to phenomenology with an appeal to stimulus-control. The precise verdicts of this border-marker would thus depend on the reach of stimulus-controlled phenomenology.

Similar remarks apply to episodic memories. As far as phenomenology goes, introspection seems to put them in the same category as paradigmatically perceptual experiences. For example, there doesn't seem to be any difference in kind between what it's like to hear one's favorite song played live versus what it's like to replay that auditory experience the next day.<sup>19</sup>

But once again, the functions performed by episodic memories preclude them from being typed as (narrowly) perceptual states: they do not have a constitutive role to play in solving instances of the underdetermination problem.<sup>20</sup>

### *7.1.1. Denying the Same-Phenomenology Thesis*

Despite its initial plausibility, there is room to doubt the same-phenomenology thesis. Some disjunctivists claim that the phenomenology of visualizing or hallucinating an object is different, in kind, to the phenomenology of actually seeing one (see Martin 2002). Many find claims of this sort obscure. But, whatever its merits, the disjunctivist gambit just brings us back to the necessity of stimulus-control. I can hardly see or hear an object without the relevant processes being causally controlled by the proximal array that the object produces: that is precisely what distinguishes perceptions from visualizations and hallucinations. Thus, if we were to follow the disjunctivist and mark a perception/cognition border in terms of the

perceptual relation, this would be tantamount to marking that border in terms of whatever species of stimulus-control happens to be constitutive of that relation.

There are ways to reject the same-phenomenology thesis without being a disjunctivist. However, an appeal to stimulus-control still seems unavoidable. For instance, Kriegel (2015a, 2015b) argues that, in contrast to genuine perceptions, episodic memories present states of affairs *as past*, while top-down forms of imagery present states of affairs *as unreal*. Kriegel claims that these phenomenal aspects show up in states' attitudinal components, not in their contents. But what makes this the case? It is very hard to answer this question without invoking stimulus-control. For example, the fact that my visualization presents a landscape *as unreal*, whereas, my current visual perception presents the surrounding landscape *as real*, is plausibly grounded in the fact that the latter functions so as to be causally controlled by proximal stimuli, while the former does not.

## 7.2. Content

I have just argued that we cannot mark a border between perception and cognition by invoking differences in phenomenology alone: stimulus-control must be brought into the picture. Similar remarks apply to the claim that the perception/cognition border is marked by a difference in content-type.

The thesis that perceptual states do not possess the same type of content as cognitive ones has appeared in several different guises over the past few decades. In the hands of theorists such as Evans (1982), Crane (1992, 2009), and Peacocke (2001), the claim is that perceptual states are distinguished by their nonconceptual contents. For instance, Evans (1982, 122–9, 154–60) argues that the contents of perception are more fine-grained than those of cognition, where the best explanation for this difference involves construing the former as nonconceptual. Evans (1982, Ch 4) also argues that perceptual states fail to meet the *Generality Constraint*, which says that if a subject undergoes a mental state with the conceptual content that *a* is *F*, she must have the resources to represent *a*'s being *G*, where *G* is some property of which she also has a concept. Those who see perceptual states as failing to meet the *Generality Constraint* often construe them as both nonconceptual and lacking in propositional structure (see Burge 2010; Crane 1992). Importantly, this lack of propositional structure is seen as delimiting the types of transitions in which perceptual states can participate. For instance, Crane (1992, 17) claims that there are no deductive inferences between perceptions.

In the hands of theorists such as Dretske (1981, Ch 6) and Kulvicki (2004, 2007, 2015), the different-content thesis becomes the claim that the contents of perception are analog, whereas the contents of cognition are digital: a representation is said to carry the information that *a* is *F* in digital form just in case *being F* is the only information it carries about *a*, whereas a representation of the same state of affairs is analog just in case it always carries further information about *a*. For example, while I can think that the banana I'm eating is yellow without my thought also representing its shape and position, I can't seem to visually perceive the banana's color without also perceiving its shape and position.

I won't dip my toes into the vast debate concerning the existence of nonconceptual content. My goal in what follows is just to show that however one construes the contents of paradigmatically perceptual states, there are compelling reasons to construe certain paradigmatically cognitive states as possessing the same type of content.

If we trust introspection as a reliable source of evidence concerning differences in content-type then one of the arguments for construing voluntary imagery as having the same type of content as perception is a familiar one: namely, subjects often fail to distinguish between those aspects of their overall experience's content that are contributed by imagination, and those that are contributed by veridical perception. A natural explanation of subjects' inability to distinguish these two aspects is that they possess the same type of content.

No doubt, we often succeed in distinguishing mental images from genuine perceptions by accessing content-related differences. For example, I might recognize the target state as a *visualization* of a banana because its content is, in some sense, less rich than the genuine perception of one would normally be—more specifically, perhaps my visualization represents the banana as yellow, without representing which determinate of yellow it instantiates. But cases such as this do not undermine the thesis that visualizations and genuine perceptions have the same *type* of content, for the content-difference invoked above is a difference in degree, not kind.<sup>21</sup>

Additional support for the same-content thesis comes from a consideration of those interactions that occur between processes of imagination and paradigmatically perceptual processes. Berger and Ehrsson (2013) recently carried out a series of experiments in which classic cross-modal illusions were altered in such a way that one of the contributing processes stemmed from voluntary imagination. For example, in their variant of the cross-bounce illusion, subjects were asked to imagine a sound while watching two objects meet. As in the original illusion, the objects were experienced as bouncing off, as opposed to passing through, one another. Berger and Ehrsson (2013) also demonstrated that a variant of the McGurk illusion can be induced by getting subjects to imagine a “ba” sound while watching the lip movements corresponding to a “ga” sound. As with the original illusion, subjects ended up experiencing a “da” sound.

If we adopt the same-content thesis, the existence of imagination-induced variants of cross-modal illusions makes sense. It is well known that imagining a stimulus recruits the same neural structures as perceiving one (see Kosslyn 2005). There are also behavioral resemblances. For instance, eye movements made during episodes of mental imagery mimic those that occur during the perception of the relevant type of scene (Laeng *et al.* 2014). Subjects also have difficulty imagining a scene when asked to keep their eyes still during visualization (Laeng & Teodorescu 2002). Finally, the dilation of the pupil reflects the darkness/brightness of the imagined scene (Laeng & Sulutvedt 2014).<sup>22</sup> If we pair these findings with the assumption that imagining a given stimulus produces a state with the same type of content as its perceptual counterpart, imagination-induced versions of cross-modal illusions are to be expected: they are cases in which low-level sensory systems fail to distinguish between inputs from other sensory systems, and inputs (with like-contents) from imagination.

On the other hand, if there were a difference in type between the contents of imagination and the contents of paradigmatic perception, we wouldn't predict cross-modal influences to persist when stimulus-driven influences are swapped out for imagination-based ones. Cross-modal effects occur because of the mechanisms that weigh the reliability of competing sources of information—for example, in general, vision provides the most reliable spatial information, whereas audition provides the most reliable temporal information. Cross-modal *illusions* occur when weights are not apportioned in an optimal fashion (e.g. too much weight is given to visual information in the McGurk effect). Why would these weighting mechanisms take inputs from top-down visualization if top-down visualization does not share the same type of content as stimulus-driven sources of visual information? One of the central claims about proprietary perceptual content is that states possessing it enter into different types of transitions than states that do not—recall that according to Crane (1992), perceptual states do not have the requisite type of content to enter into deductive inferences. If this is true, why would cross-modal weighting mechanisms be fooled by inputs from top-down imagination? At the very least, rejecting the same-content thesis makes the answer to this question obscure.

The thesis that episodic memories possess the same type of content as paradigmatic perceptions is also compelling. From the first-person perspective, mental images and episodic memories certainly appear to share the same type of content. This point is bolstered by the fact that it is easy to forget whether an episode one is representing is an actual memory, or, just an image that one has formed. For example, you might be unsure whether the conversation you are 'playing' in your mind right now is an actual memory, or, just an imagined conversation in which you say the things you wished you had said. The false memories that are produced via certain forms of recovered-memory therapy (RMT)—a technique that often utilizes hypnosis and guided imagery—probably fall into this category as well (see Loftus and Davis 2006). Given that we have already construed mental imagery as possessing the same type of content as paradigmatic perceptions, we should therefore say the same about episodic memories and paradigmatic perceptions.

Indirect support for the thesis that episodic memories share the same type of content as paradigmatic perceptions also comes from a consideration of neuroanatomical resemblances between the two. For instance, neurons activated during perceptual experiences are selectively reactivated during free recall of those experiences (see Wheeler *et al.* 2000; Gelbard-Sagiv *et al.* 2008).

There are thus several reasons for thinking that some paradigmatically cognitive states possess the same type of content as paradigmatically perceptual ones. As was the case with the appeal to phenomenology, what this means is that without invoking stimulus-control, appeals to content-type alone will not allow us to mark a recognizable perception/cognition border.<sup>23</sup>

### 7.2.1. *Denying the Same-Content Thesis*

Of course, some theorists deny that perceptual states have content in the first place. Relationalists argue that instead of representing distal entities, perceptual states put us into non-representational relations with them (Brewer 2006; Travis 2009). If the

relationalists are right then the same-content thesis is false. However, for reasons that should be familiar, adopting relationalism would not undermine the thesis that any attempt to mark a border between perception and cognition must invoke stimulus-control. The relation that relationalists construe as constitutive of being a perceptual state just is the perceptual relation, and as was pointed out above, whatever that relation may be, it is constituted, in part, by stimulus-control.

### 7.3. *Format*

Finally, consider the claim that perceptual states have an iconic format, while cognitive states have a discursive one (Dretske 1981; Kulvicki 2004, 2015; Fodor 2007; Carey 2009; Raftopoulos 2009; Burge 2010, 2014; Block 2014). For the purposes of illuminating this distinction, consider the way in which a picture carries its content versus the way in which a sentence of natural language does. Every part of a picture represents part of the scene—Fodor (2007) calls this *the picture principle*—whereas not every part of a sentence does. For example, every pixel in a photograph of my cat, Oliver, represents part of the scene containing him. But the same does not hold true of the sentence “Oliver is purring”: the letter “O” does not represent anything at all. There is thus a fundamental difference in the way each vehicle carries its content.<sup>24</sup>

The iconic/discursive distinction is a syntactic distinction between two types of content-carrying vehicles. Nonetheless, claims about differences in format go hand in hand with claims about differences in content. For instance, discursive states are thought to have a distinctive compositional semantics, such that the meaning of the whole is exhaustively determined by the meanings of its atomic constituents along with its syntax. In contrast, even though iconic states have their own compositional semantics, the meaning of the whole is not determined by the meanings of any atomic constituents: they simply do not have canonical decompositions like discursive states do. Beyond this, characterizations of the iconic/discursive distinction are controversial. For instance, some theorists claim that the picture principle entails that iconic states possess nonconceptual contents (see Fodor 2007).

For our purposes, the most pressing question is this: why think that the iconic/discursive distinction maps onto the perception/cognition distinction? Suppose you take yourself to have independent reasons for thinking that perceptual content is proprietary: e.g. it is fine-grained and lacks propositional structure. In that case, construing perceptual states as iconic provides a way of explaining this fact.

It is beyond the scope of this paper to advance the debate about whether paradigmatically perceptual states are iconic, and whether this imbues them with nonconceptual content. As was the case above, my goal is just to defend the claim that invoking differences in format alone will not give us a recognizable perception/cognition border.

Given that differences in content-type are almost always explained in terms of differences in format-type, arguments for the view that visualizations, episodic memories, and paradigmatically perceptual states all share the same type of format are the same ones that are used to argue for the same-content thesis. For instance, if we couple the fact that visualizations recruit the same neural structures as

paradigmatic perceptions with the assumption that they also share the same type of format, we would predict imagery-induced variants of cross-modal illusions. This is just not a phenomenon you would predict if you reject the same-format thesis.

I won't rehearse the rest of the arguments here. The point is that if you adopt the same-content thesis, you should adopt the same-format thesis. And if you adopt the same-format thesis then invoking differences in format-type will not be sufficient to mark a recognizable perception/cognition border—stimulus-control must be brought into the picture.<sup>25</sup>

## 8. Adaptation and Constancy

In addition to traditional border-markers, two others have been discussed in the recent literature: selective adaptation and constancy. Selective adaptation is the process by which a neuron's activity decreases amidst prolonged exposure to its preferred stimulus. One effect of this decrease in activity is that it raises the threshold for detecting the stimulus in question, thereby biasing the percept toward alternative stimuli. For example, in the waterfall illusion, staring at the water's motion reduces the firing rate of those neurons tuned to downward motion, thereby biasing the visual system in favor of representing upward motion. This means that if one views a stationary scene immediately after having stared at a waterfall, one will have an illusory experience of upward motion—this is referred to as a *motion aftereffect*.

Adaptation is ubiquitous when it comes to paradigmatically perceptual processes, such as color perception; volume perception; heat perception; and so on. In fact, adaptation is so ubiquitous when it comes to paradigmatically perceptual phenomena that Block (2014) has put it forward as a marker of the perceptual: one we can use to sort out the controversial cases.

In assessing adaptation's candidacy as a border-marker, consider the fact that voluntary imagery can produce adaptation effects. For instance, Winawer *et al.* (2010) induced a motion aftereffect by asking subjects to simply *visualize* motion, before exposing them to an array of stationary dots. If upward motion was visualized, the stationary dots appeared to move downward; but if downward motion was visualized, the same dots appeared to move upward. The converse has also been demonstrated. For instance, Gildea *et al.* (1995) found that viewing motion induces an aftereffect at the level of imagination.

As was explained above, voluntary visualizations are broadly cognitive states. Thus, given that they exhibit adaptation effects, the appeal to adaptation will only permit us to distinguish broad cognition from narrow perception if we couple it with an appeal to stimulus-control. On this way of marking a perception/cognition border, perceptual processes are those that exhibit stimulus-controlled-adaptation.

Similar remarks apply to constancy. Constancy mechanisms are what enable us to perceive the same distal feature across variations in the proximal stimulus. For instance, I see you as having the same size regardless of whether you are standing right in front of me or far off in the distance. Like adaptation, constancy is ubiquitous when it comes to paradigmatically perceptual processes, leading some theorists—most notably, Burge (2010)—to deploy it as a marker of the perceptual.

Are there cognitive constancies? Burge certainly thinks of perceptual constancies as analogues of Fregean informative identity statements (2010, 411). However, in discussing Frege cases, he stops short of using the term ‘cognitive constancies,’ and with good reason, for we must not elide the distinction between the following: (1) distinct thoughts that pick out the same object via co-referring terms; and (2) distinct representations that pick out the same distal feature despite variations in the proximal stimuli that causally control these representations. It is the latter that perceptual psychologists are referring to when they deploy the theoretical term, ‘constancy.’

## 9. Pluralism Versus Eliminativism

Our border search has uncovered more than one viable way of marking a border between perception and cognition. The natural conclusion to draw is that we should be pluralists. According to the pluralist, there are multiple ways of marking a border between perception and cognition, each one being tethered to a distinct explanatory goal. But if the borders we have uncovered are each motivated by a distinct explanatory goal, why not join the eliminativists and jettison the notions of perception and cognition altogether?

### 9.1. *The Class of Stimulus-Controlled States as a Superordinate Natural Kind*

Recall that according to some eliminativists, the ubiquity of top-down influences undermines the explanatory significance of the perception/cognition distinction. Given that there are explanatorily fruitful ways of drawing the distinction that do not require states on the perceptual side of the divide to be encapsulated, this route to eliminativism is blocked. The more promising route takes its inspiration from eliminativist arguments in the philosophy of biology—in particular, those surrounding the explanatory significance of species.

It is widely held that several different species categories find a home in biology, each one being deployed in the service of a different explanatory project (Kitcher 1984; Dupré 1993). For example, there is the biological species concept, which according to Ernst Mayr pick out “groups of interbreeding natural populations that are reproductively isolated from other such groups” (Mayr 1969: 26); ecological niche concepts, which demarcate species in terms of the ecological niches they occupy; and several others.

In reacting to this proliferation of categories, one reaction is to go pluralist. But another reaction is to conclude that the species concept has outlived its usefulness in biology. One theorist who adopts this second line is Ereshefsky (1998). According to Ereshefsky, the various species categories on offer do not share any distinctive commonality.

What do these different types of species have in common that renders them species? If species taxa lack a common unifying feature, then we have reason to doubt the existence of the species category. (1998, 111)

A disjunctive definition of the species category would not tell us why various taxa are species . . . disjunctive definitions lack ontological import (1998, 115).

As Devitt (2011) points out, Ereshefsky's inference from the claim that a category is disjunctive to the claim that that category does not exist is too quick. Disjunctive classes arguably exist: it is just that they are theoretically uninteresting.

Nonetheless, Ereshefsky's question about unifying features is a good one. If those states that are rendered perceptual by at least one of the border-markers uncovered above form a heterogeneous class of unrelated kinds, shouldn't we conclude that the notion of the perceptual has no place in cognitive science (*mutatis mutandis* for the notion of the cognitive)? If we were to endorse this line of thinking, perceptual and cognitive categories would thus be in the same basket as other notions that have no explanatory significance in cognitive science, such as those mental processes taking more than 80 ms to occur, and those desires tokened by those wearing their hats backwards.

By now, it should be obvious which feature I see as unifying the class of perceptual categories into an explanatorily significant kind. The distinctive commonality in question is stimulus-control. Recall that in carrying out our border search, each of the explanatorily significant border-markers that we uncovered classifies a state as perceptual only if it exhibits stimulus-control: the most liberal classificatory scheme deploying the unqualified notion of stimulus-control; the other schemes deploying some qualified form, such as stimulus-specific-control or encapsulated-stimulus-control.

The different perceptual categories that we have catalogued are therefore members of a superordinate natural kind (*mutatis mutandis* for the complimentary set containing the various categories of the cognitive). Examples of superordinate kinds with explanatory significance abound in science. The class of *acids*, whose subkinds include the class of substances with a pH of 4 as well as the class of substances with a pH of 5, is one such example; as is the class of objects with mass, whose subkinds include the class of 5kg-objects and the class of 2kg-objects.

It would be implausible to deny that these superordinate kinds earn their explanatory keep in the relevant sciences. Thus, to deny that the class of perceptual categories earns its explanatory keep in cognitive science just because it is a superordinate class of explanatorily significant subkinds would be to ignore the pattern that we see across the sciences—hierarchies of explanatory significant kinds are the rule, not the exception.<sup>26</sup>

### 9.2. *Pluralism Versus Indeterminacy*

Couldn't the monist recognize the various border-candidates uncovered above, and declare that it is indeterminate which one the perception/cognition distinction refers to?<sup>27</sup>

As I see it, the question of when to go pluralist and when to posit indeterminacy depends on whether the candidates in question each fulfill the same explanatory role. For instance, in claiming that "rabbit" is indeterminate between the set of rabbits and the set of undetached-rabbit-parts, Quine (1960) is claiming that we can

explain verbal behavior equally well, regardless of which reference scheme we go with.

Pluralism becomes the preferred option when the candidates in questions fulfill different—even if partially overlapping—explanatory roles. Given that the various border-markers we have uncovered serve distinct explanatory roles, we should favor pluralism over monism-plus-indeterminacy.

### 10. Ramifications for Ongoing Debates

So we should be pluralists, not eliminativists or monists, about the perception/cognition distinction. In concluding, I briefly outline implications for ongoing debates in philosophy of mind.

Take the vexed debate concerning which categories are represented in perception. According to conservatives, perception only represents low-level categories, such as color, shape, and motion (see Tye 1995; Clark 2000). According to liberals, perception also represents various high-level categories, such as *being a tiger* (see Peacocke 1992; Siegel 2010).

There is a striking lack of consensus concerning how to adjudicate this debate. Bayne (2009) flirts with the idea that the debate is terminological. Similarly, in the absence of a clear distinction between perception and cognition, Lycan (2014) suggests that the question of whether perception represents high-level categories may be a bad one. Pluralism provides a clear way of fleshing out these nascent concerns. The debate is intractable because questions about what perception represents only make sense once we have made it clear which perception/cognition border we are invoking. If one's explanatory project requires an appeal to the unqualified notion of stimulus-control, the high-level categories that are represented via processes of rapid object recognition will be counted among the contents of perception. However, if one's explanatory project is centered on the distinction between those stimulus-controlled states that are nonconceptual, and those that are conceptual, these same categories will be counted among the cognitively represented.

Similarly, there is much confusion in the literature as to what counts as an instance of cognitive penetration. As long as parties to the debate have neglected to make the explanatory context clear, this confusion is unsurprising. Questions about whether a top-down process qualifies as a cognitive one can only be answered once the explanatory context—and thereby the perception/cognition border that it engenders—has been made explicit.

Getting clear on the explananda that motivate pluralism about the perception/cognition distinction will therefore do more than just reveal the mind's joints—a worthy pursuit in its own right. It also promises to shed light on core debates in philosophy of mind. Given the widespread failure to heed pluralism, it is not surprising that these debates have become vexed.<sup>28</sup>

### Notes

<sup>1</sup> Some recent monists include Raftopoulos (2009); Burge (2010, 2014); Block (2014); and Firestone & Scholl (2016).

<sup>2</sup> See Jackson (1998, 30–31) and Ramsey (2007, 13).

<sup>3</sup> See Stich (1996) for an extended discussion.

<sup>4</sup> *Broad* perception plausibly constitutes a natural kind as well. Natural kinds can overlap: the class of chordates and the class of mammals are two such kinds.

<sup>5</sup> Similar points are made by Beck (forthcoming), Lyons (2015), and Block (2016, 21–22).

<sup>6</sup> See Schellenberg (2016) for a review.

<sup>7</sup> Transduction is the process whereby one form of energy is transformed into another. For example, in the case of vision, rods and cones in the retina transform light energy into electrochemical energy, which is transmitted to the brain via the optic nerve.

<sup>8</sup> See Godfrey-Smith (1994).

<sup>9</sup> See Bigelow and Pargetter (1987). Other ahistorical accounts include Cummins (1975), and Nanay (2010).

<sup>10</sup> At one point, Burge (2010, 378) characterizes perceptual states as those that are under the “causal control” of the stimulus. However, he does not elaborate on what he means by this term. According to Beck (forthcoming), perceptual states are “causally sustained” by present proximal stimulation, whereas, cognitive states are not.

<sup>11</sup> In what follows, I’ll employ the neutral term ‘stimulus-controlled’ to refer to the class of states that perform the function described above. I’ll reserve the term ‘*Stimulus-Control<sub>P/C</sub>*’ for the stronger claim that the class of stimulus-controlled states constitutes the class of perceptual states.

<sup>12</sup> The underdetermination problem is also referred to as *the inverse problem*. See Palmer (1999, 23).

<sup>13</sup> I’ll say that a state is causally controlled in a ‘stimulus-specific’ manner if it performs the function described above. I’ll reserve the term ‘*Stimulus-Specific-Control<sub>P/C</sub>*’ for the thesis that this makes the state in question perceptual, not cognitive.

<sup>14</sup> Thanks to an anonymous referee for raising this issue.

<sup>15</sup> Vision scientists recognize at least two systems for the brief retention of visual information: iconic memory and VSTM. Iconic memory has a much higher capacity than VSTM; however, it decays more rapidly (see Matsukura & Hollingworth 2011).

<sup>16</sup> See Beck (forthcoming) for a slightly different take on demonstrative thoughts.

<sup>17</sup> In fact, there is also physiological evidence for stimuli-non-specific representations of relatively low-level features, such as motion (Konkle *et al.* 2009) and three-dimensional shape (Amedi *et al.* 2001). These representations are thus good candidates for states that satisfy *Stimulus-Control<sub>P/C</sub>*, but not *Stimuli-Specific-Control<sub>P/C</sub>*.

<sup>18</sup> See Firestone and Scholl (2016) for a critical review of purported demonstrations of cognitive penetration—they count at least 160 in the literature.

<sup>19</sup> Of course, there is usually a difference in richness or vivacity. See Nanay (2015) for a helpful discussion of this issue.

<sup>20</sup> Possessing a certain type of phenomenology may well be what unites the class of *broadly* perceptual states, distinguishing them from the class of *narrowly* cognitive ones. However, as I explained above, my concern is with the narrow-perception/broad-cognition border, and the point I’m making above is that a bare appeal to differences in phenomenology will not enable us to mark *this* border.

<sup>21</sup> There are various ways to cash out the claim that the contents of perception and imagination differ in degree, but not kind. Take Kulvicki’s (2007) thesis that perceptual content is “vertically articulate,” whereas the contents of thought are not. A state has vertically articulate content just in case in representing *P*, it also represents some *Q*, such that the latter is an abstraction from *P* (2007, 359). For instance, whereas my thought that the scarf is scarlet does not (necessarily) represent the scarf as red and as colored, my perception of the scarf does represent these less determinate properties (i.e. these abstractions). If the contents of perception and imagination differ in degree, not kind, one way to explicate the difference is via an appeal to degrees of abstraction. For instance, perhaps my visualization of the scarf is less determinate, and involves fewer abstractions than my perception of it.

<sup>22</sup> I’m indebted to Bence Nanay (2015) for these examples.

<sup>23</sup> Nanay (2015) argues the contents of perception are (normally) more determinate than the contents of imagination because the former come in a stimulus-driven manner from what you perceive, whereas the latter are formed in a stimulus-independent manner. Nanay’s version of the same-content thesis is thus very much in the spirit of my core thesis that any marker of the perceptual must appeal to the notion of stimulus-control.

<sup>24</sup>The different-format thesis is sometimes cashed out as the view that perceptual states have a map-like format, whereas cognitive states have a discursive one (see Rescorla 2009 for a discussion).

<sup>25</sup>I do not have the space to uncover points of divergence between *Stimulus-Control<sub>P/C</sub>*, *Stimuli-Specific-Control<sub>P/C</sub>*, and the potential border-markers discussed in this section. In brief, though, one obvious point of divergence concerns object recognition. It is widely held that object recognition is a conceptual capacity, and yet, as I argued in section 5, rapid recognition plausibly satisfies *Stimulus-Control<sub>P/C</sub>*. Rapid recognition is therefore classed as perceptual if we go with *Stimulus-Control<sub>P/C</sub>*, but cognitive if we construe the perceptual as the *stimulus-controlled-cum-non-conceptual*.

<sup>26</sup>See Weiskopf (2009) who makes the same type of point in favor of pluralism about concepts. See also Phillips (2016).

<sup>27</sup>Thanks to Ned Block for raising this question.

<sup>28</sup>For helpful feedback, I would like to thank Jacob Beck, Ned Block, Uriah Kriegel, Bence Nanay, Nicholas Shea, and an anonymous referee. I would also like to thank John and Kathy Larocca for providing a peaceful place to write.

## References

- Amedi, A., Jacobson, G., Hendler, T., Malach, R., and Zohary, E. (2002). Convergence of Visual and Tactile Shape Processing in the Human Lateral Occipital Complex. *Cerebral Cortex*, 12(11), 1202–1212.
- Auvray, M. and Spence, C. (2008). The Multisensory Perception of Flavor. *Consciousness and Cognition*, 17(3), 1016–1031.
- Bayne, T. (2009). Perception and the Reach of Phenomenal Content. *The Philosophical Quarterly*, 59(236), 385–404.
- Bayne, T. (2014). The Multisensory Nature of Perceptual Consciousness. In D. Bennett and C. Hill (eds.), *Sensory Integration and the Unity of Consciousness*. MIT Press (pp. 15–36).
- Bayne, T. and Montague, M. (2011). *Cognitive Phenomenology*. Oxford: Oxford University Press.
- Beck, J. (forthcoming). Marking the Perception-Cognition Boundary: The Criterion of Stimulus-Dependence. *Australasian Journal of Philosophy*.
- Berger, C. and Ehrsson, H. (2013). Mental Imagery Changes Multisensory Perception. *Current Biology*, 23, 1367–1372.
- Bigelow, J. and Pargetter, R. (1987). Functions. *Journal of Philosophy*, 84, 181–197.
- Block, N. (2014). Seeing-As in the Light of Vision Science. *Philosophy and Phenomenological Research*, 89(1), 560–572.
- Block, N. (2016). Tweaking the Concepts of Perception and Cognition. *Behavioral and Brain Sciences*, Commentary on Firestone & Scholl (2016), 21–22.
- Brewer, B. (2006). Perception and Content. *European Journal of Philosophy*, 14, 165–81.
- Burge, T. (2010). *Origins of Objectivity*. Oxford: Oxford University Press.
- Burge, T. (2014). Reply to Block: Adaptation and the Upper-Border of Perception, 89(3), 573–583.
- Carey, S. (2009). *The Origin of Concepts*. Oxford: Oxford University Press.
- Carruthers, P. (2006). *The Architecture of the Mind: Massive Modularity and the Flexibility of Thought*. Oxford: Clarendon Press.
- Clark, A. (2000). *A Theory of Sentience*. Oxford: Oxford University Press.
- Clark, A. (2013). Whatever next? Predictive Brains, Situated Agents, and the Future of Cognitive Science. *Behavioral and Brain Sciences*, <https://doi.org/10.1017/S0140525X12000477>
- Cosmides, L. and Tooby, J. (1992). *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. New York: Oxford University Press.
- Crane, T. (1992). The Nonconceptual Content of Experience. In T. Crane (ed.), *The Contents of Experience*. Cambridge: Cambridge University Press (pp. 136–157).
- Crane, T. (2009). Is Perception a Propositional Attitude? *The Philosophical Quarterly*, 59(236), 452–469.
- Cummins, R. (1975). Functional Analysis. *Journal of Philosophy*, 72, 741–765.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Oxford: Blackwell.
- Dupré, J. (1993). *The Disorder of Things*. Cambridge, MA.: Harvard University Press.
- Ereshefsky, M. (1998). Species Pluralism and Anti-Realism. *Philosophy of Science*, 65, 103–120.

- Evans, G. (1982). *The Varieties of Reference*. Oxford: Oxford University Press.
- Firestone, C. and Scholl, B. (2016). Cognition Does Not Affect Perception: Evaluating the Evidence for ‘Top-Down’ Effects. *Behavioral and Brain Sciences*, doi: <http://doi.org/10.1017/S0140525X15000965>
- Fodor, J. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Fodor, J. (2007). Revenge of the Given. In B. McLaughlin & J. Cohen (eds.), *Contemporary Debates in Philosophy of Mind*. Oxford: Blackwell (pp. 105–116).
- Gilden, D., Blake, R., and Hurst, G. (1995). Neural Adaptation of Imaginary Visual Motion. *Cognitive Psychology*, 28(1), 1–16.
- Gelbard-Sagiv, H., Mukamel, R., Harel, M., Malach, R., and Fried, I. (2008). Internally Generated Re-activation of Single Neurons in Human Hippocampus During Free Recall. *Science*, 322(5898), 96–101.
- Godfrey-Smith, P. (1994). A Modern History Theory of Functions. *Noûs*, 28, 344–362.
- Godfrey-Smith, P. (2000). On the Theoretical Role of Genetic Coding. *Philosophy of Science*, 67, 26–44.
- Grill-Spector, K. and Kanwisher, N. (2005). Visual Recognition: As Soon as You Know It Is There, You Know What It Is. *Psychological Science*, 16(2), 152–160.
- Hollingworth, A., Richard, A., and Luck, S. (2008). Understanding the Function of Visual Short-Term Memory: Transsaccadic Memory, Object-Correspondence, and Gaze Correction. *J Exp Psychol Gen.*, 137(1), 163–181.
- Jackson, F. (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
- Jordan, K., Clark, K., and Mitroff, S. (2010). See an Object, Hear an Object File: Object Correspondence Transcends Sensory Modality. *Visual Cognition*, 18(4), 492–503.
- Kafaligonul, H., Breitmeyer, B., and Ogmen, H. (2015). Feedforward and Feedback Processes in Vision. *Frontiers in Psychology*, DOI: <http://doi.org/10.3389/fpsyg.2015.00279>
- Kahneman, D., Treisman, A., and Gibbs, B. (1992). The Reviewing of Object Files: Object-Specific Integration of Information. *Cognitive Psychology*, 24, 175–219.
- Kanwisher, N. (2001). Neural Events and Perceptual Awareness. *Cognition*, 79, 89–113.
- Keyesers, C., Xiao, D., Földiák, P., and Perrett, D. (2001). The Speed of Sight. *Journal of Cognitive Neuroscience*, 13(1), 90–101.
- Kitcher, P. (1984). Species. *Philosophy of Science*, 51, 308–333.
- Konkle, T., Wang, Q., Hayward, V., and Moore, C. (2009). Motion Aftereffects Transfer Between Touch and Vision. *Current Biology*, 19, 745–750.
- Kosslyn, S. (2005). Mental Images and the Brain. *Cognitive Neuropsychology*, 22(3/4), 333–347.
- Kriegel, U. (2015a). *The Varieties of Consciousness*. Oxford: Oxford University Press.
- Kriegel, U. (2015b). Experiencing the Present. *Analysis*, 75, 407–413.
- Kulvicki, J. (2004). Isomorphism in Information-Carrying Systems. *Pacific Philosophical Quarterly*, 85(4), 380–395.
- Kulvicki, J. (2007). Perceptual Content is Vertically Articulate. *American Philosophical Quarterly*, 44(4), 357–369.
- Kulvicki, J. (2015). Analog Representation and the Parts Principle. *Review of Philosophy and Psychology*, 6, 165–180.
- Laeng, B., Bloem, I., D’Ascenzo, S., and Tommasi, L. (2014). Scrutinizing Visual Images: The Role of Gaze in Mental Imagery and Memory. *Cognition*, 131, 263–283.
- Laeng, B. and Sulutvedt, U. (2014). The Eye Pupil Adjusts to Imaginary Light. *Psychological Science*, 25(1), 188–197.
- Laeng, B. and Teodorescu, D. (2002). Eye Scanpaths During Visual Imagery Re-enact Those of Perception of the Same Visual Scene. *Cognitive Science*, 26, 207–231.
- Loftus, E. and Davis, D. (2006). Recovered Memories. *Annu. Rev. Clin. Psychol.*, 2, 469–498.
- Lupyan, G. (2015a). Cognitive Penetrability of Perception in the Age of Prediction: Predictive Systems are Penetrable Systems. *Review of Philosophy and Psychology*, 6(4), 547–569.
- Lupyan, G. (2015b). Reply to Macpherson: Further Illustrations of the Cognitive Penetrability of Perception. *Review of Philosophy and Psychology*, 6(4), 585–589.

- Lycan, W. (2014). What Does Vision Represent? In B. Brogaard (ed.), *Does Perception Have Content?* Oxford: Oxford University Press (pp. 311–328).
- Lyons, J. (2015). Unencapsulated Modules and Perceptual Judgment. In J. Zeimbekis and A. Raftopoulos (eds.), *The Cognitive Penetrability of Perception: New Philosophical Perspectives*. Oxford: Oxford University Press (pp. 103–122).
- Man, K., Kaplan, J., Damasio, A., and Meyer, K. (2012). Sight and Sound Converge to Form Modality-Invariant Representations in Temporoparietal Cortex. *Journal of Neuroscience*, 32(47), 16629–16636.
- Mandelbaum, E. (forthcoming). Seeing and Conceptualizing: Modularity and the Shallow Contents of Perception. *Philosophy and Phenomenological Research*, <https://doi.org/10.1111/phpr.12368>
- Martin, M. (2002). The Transparency of Experience. *Mind & Language*, 17, 376–425.
- Matsukura, M. and Hollingworth, A. (2011). Does Visual Short-Term Memory Have a High Capacity Stage? *Psychon Bull Review*, 18(6), 1098–1104.
- Mayr, E. (1969). *Principles of Systematic Zoology*. New York: McGraw Hill.
- McFarland, D. and Cacace, A. (1999). Defining Perception and Cognition. *Behavioral and Brain Sciences*, 22(3), 383–385.
- McGurk, H. and MacDonald, J. (1976). Hearing Lips and Seeing Voices. *Nature*, 264, 746–748.
- Nanay, B. (2010). A Modal Theory of Function. *Journal of Philosophy*, 107, 412–431.
- Nanay, B. (2015). Perceptual Content and the Content of Mental Imagery. *Philosophical Studies*, 172, 1723–36.
- Nudds, M. (2015). Is Audiovisual Perception ‘Amodal’ or ‘Crossmodal’? In D. Stokes, M. Matthen, and S. Biggs (eds.), *Perception and Its Modalities*. Oxford: Oxford University Press (pp. 166–190).
- Palmer, S. (1999). *Vision Science: Photons to Phenomenology*. Cambridge, MA: MIT Press.
- Peacocke, C. (1992). *A Study of Concepts*. Cambridge, MA.: MIT Press.
- Peacocke, C. (2001). Does Perception Have a Nonconceptual Content? *Journal of Philosophy*, 98, 239–64.
- Perky, C. (1910). An Experimental Study of Imagination. *The American Journal of Psychology*, 21(3), 422–452.
- Phillips, B. (2016). Contextualism About Object-Seeing. *Philosophical Studies*, 173, 2377–2396.
- Potter, M., Wyble, B., Haggmann, C., and McCourt, E. (2014). Detecting Meaning in RSVP at 13 ms Per Picture. *Attention, Perception, & Psychophysics*, 76(2), 270–279.
- Pylyshyn, Z. W. (1984). *Computation and Cognition*. Cambridge, MA: MIT Press.
- Pylyshyn, Z. W. (2003). *Seeing and Visualizing: It's Not What You Think*. Cambridge, MA: MIT Press.
- Quine, W. V. O. (1960). *Word and Object*. Cambridge, MA.: MIT Press.
- Quiroga, R., Kraskov, A., Koch, C., and Fried, I. (2009). Explicit Encoding of Multimodal Percepts by Single Neurons in the Human Brain. *Current Biology*, 19(15), 1308–1313.
- Raftopoulos, A. (2009). *Cognition and Perception: How Do Psychology and Neural Science Inform Philosophy?* MIT Press.
- Ramsey, W. (2007). *Representation Reconsidered*. Cambridge, MA.: Cambridge University Press.
- Rescorla, M. (2009). Cognitive Maps and the Language of Thought. *British Journal for the Philosophy of Science*, 60(2), 377–407.
- Schellenberg, S. (2016). Perceptual Particularity. *Philosophy and Phenomenological Research*, 93(1), 25–54.
- Shams, L. and Kim, R. (2010). Crossmodal Influences on Visual Perception. *Physics of Life Reviews*, 7(3), 269–284.
- Shea, N. (2015). Distinguishing Top-Down from Bottom-Up Effects. In D. Stokes, M. Matthen, and S. Biggs (eds.), *Perception and Its Modalities*. Oxford: Oxford University Press (pp. 73–91).
- Siegel, S. (2010). *The Contents of Visual Experience*. Oxford: Oxford University Press.
- Stegman, U. (2014). Causal Control and Genetic Causation. *Noûs*, 48(3), 450–465.
- Stich, S. (1996). *Deconstructing the Mind*. Oxford: Oxford University Press.
- Travis, C. (2015). ‘The Silence of the Senses’, *Mind*, 113: 57–94.
- Tye, M. (1995). *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.
- Weiskopf, D. (2009). The Plurality of Concepts. *Synthese*, 169, 145–173.
- Wheeler, M., Peterson, S., and Buckner, R. (2000). Memory’s Echo: Vivid Remembering Reactivates Sensory-Specific Cortex. *PNAS*, 97(20), 11125–11129.

- Winawer, J., Huk, A., and Boroditsky, L. (2010). A Motion Aftereffect from Visual Imagery of Motion. *Cognition*, 114, 276–284.
- Yildirim, I. and Jacobs, R. (2012). A Rational Analysis of the Acquisition of Multisensory Representations. *Cognitive Science*, 36, 305–332.
- Zeimbekis, J. and Raftopoulos, A. (2015). *The Cognitive Penetrability of Perception: New Philosophical Perspectives*. Oxford: Oxford University Press.