Some Hallucination is Experience of the Past

Michael Barkasi

May 28, 2020

Abstract

When you hallucinate an object you are not in the normal sort of concurrent causal-sensory interaction with that object. It's standardly further inferred that the hallucinated object doesn't actually exist. But the lack of normal concurrent causal-sensory interaction does not imply that there does not exist an object that is hallucinated. It might be a past-perceived object. In this paper I argue this claim holds for at least some interesting cases of hallucination. Hallucinations generated by misleading cues (e.g., "seeing" Kanizsa triangles), hallucinations of Charles Bonnet Syndrome patients, and dreams are experience of past-perceived objects.

1 Introduction

If I visually hallucinate a pencil, the hallucinated pencil is not "really there". It's thus standardly inferred that there does not exist a physical object which is hallucinated (e.g., Macpherson 2013). In other words, there is nothing out in the world which I experience (or am in experiential contact with) as I hallucinate. But while the hallucinated pencil isn't something now in front of me with which I concurrently interact through my visual system in the normal way, it doesn't follow that there does not exist a particular pencil I'm visually experiencing. It might be that the hallucinated pencil is an object in the past: a past-viewed pencil.

In this paper I argue that this holds for at least some cases. The argument goes as follows:

P1: At least some hallucinations are generated by sensory activities which are (in part) interactions with past-perceived objects.

P2: If an experience is generated by a sensory interaction with some object O and seems to be an experience of an object O*, then O* is O.

Conclusion: At least some hallucinations are experiences of past-perceived objects.[1]

Premise 2 will take some stage setting, including explaining what I mean by 'sensory activity' and 'interaction', but with this setup should be readily accepted. This I'll do in section 2. Note also that premise 2 has an unstated *ceteris paribus* clause, but, as I'll explain, hallucinations don't raise any issues related to this clause.

Premise 1 will require more argument. The argument for it runs as follows:

P3: A sensory activity of an organism is an interaction with an object if (1) that object is a causal antecedent of the neural signals mediating the activity, and (2) the use or function to which the organism puts those signals involves that object in the right way.

P4: At least some hallucinations (and probably all normal perceptual experiences) are generated (in part) by sensory activities which are mediated by neural signals which have past-perceived objects as causal antecedents.

P5: In at least some of these hallucinations, the generating sensory activities involve those past-perceived objects in the right way.

Conclusion (P1): At least some hallucinations are generated by sensory activities which are (in part) interactions with past-perceived objects.

I will fill out this argument in section 3. In brief, premise 3 is motivated by paradigm examples of sensory interactions. Premise 4 follows from how patterns of neural activity are shaped by synaptic plasticity. An exception here will turn out to be the experiences of a fully envatted brain which has always and only ever been fed input from virtual models. For premise 5, following work by Albright (2012), what I will argue is that the sensory activity generating at least some cases of hallucination is *recall* and is similar to the sensory activity behind episodic memory. The main case is hallucinations induced by misleading cues, such as "seeing" Kanizsa triangles made of subjective contours.

More tentatively, in section 4 I suggest that two other cases of hallucination satisfy premises 4 and 5, and so might be examples of hallucinations which are experience of past objects. These cases are dreams and the clinical

case of Charles Bonnet Syndrome (CBS). Section 5 will collect together and reply to some objections to my argument.

We need not settle at the start on a rigorous definition of 'hallucination'. As a guiding concept, I take hallucinations to be states of an organism which we would intuitively characterize as experience (or *seeming* experience) of something not "really there", such that the state seems to the organism to involve the use of their senses. Prime examples include speculative cases from philosophy, such as the experiences of a life-long fully envatted brain or people trapped in the Matrix, and real-life cases like experiences induced by misleading cues, clinical cases like CBS or schizophrenia, and dreams. Hallucinations aren't mere illusions, as the latter involve experiencing an actual present distal stimulus, albeit inaccurately. I don't assume the category has sharp cutoffs, or that there is a fundamental difference between hallucinations and related cases like illusions and states of imagination. For example, "seeing" Kanizsa triangles is probably more closely related to a mere illusion like that induced by Müller-Lyer lines than it is to the experiences of an envatted brain. My argument itself will suggest that hallucinations (or some kinds of hallucinations) have much in common with, and perhaps are a kind of, imagination. Despite the ontological messiness, I take hallucinations to be unified insofar as they all involve experience (or seeming experience) of a particular object with which the subject is not concurrently interacting through their sensory systems. As noted, it's often inferred that this experienced object is (or must be) nonexistent or merely intentional, but I shall argue that in some cases the experienced object is real, albeit from the past.

When I say that some cases of hallucination involve experience of the past, what I mean is that there exists a particular object of which the hallucinating subject is aware. This existing hallucinated object is some object from the past, or, rather, a past temporal slice of an object which may or may not presently exist. To put it this way assumes an eternalist view of time on which the past exists (see Moran 2019), but if eternalism turns out false, my claim only requires minor amendment. In that case, the claim becomes that there existed some particular object of which the hallucinating subject is (now) aware.

My thesis is neutral between representationalism (e.g., Dretske 1995; Burge 2010; Siegel 2010) and views, like naïve realism (e.g., Campbell 2002; Martin 2004; Fish 2009), which reject it. There is no reason why a representationalist can't accept the claim that some hallucination is experience of past-perceived objects, perhaps saying (for example) that those hallucina-

tions are representations of past objects. I will present my argument in a way that's neutral to the debate. For example, this is why I above described hallucinations as experiences, *or seeming experiences*, of what's not really there. Representationalists standardly take them to be experiences (i.e., representations) of a nonexistent object. Naïve realists (typically) take them to be seeming experiences (i.e., seeming experiential relations) to a nonexistent object. Hereafter I'll forgo this and related terminological accommodations, leaving it to the reader to make the needed adjustments to fit their preferred view.

How does my thesis fit within broader debates in the philosophy of perception? Both (a) the commonly accepted view that sensory experience is not a relation to experienced objects (e.g., Anscombe 1965), and (b) the argument from hallucination for the claim that we are, in sensory experience, directly aware of only private inner sensibilia (e.g., Valberg 1992), are driven in large part by the assumption that hallucinated objects don't exist. If I'm correct, this assumption fails for some realistic and interesting cases of hallucination. What's left over are mostly speculative cases, such as a life-long fully envatted brain. In the final section of this paper (section 6) I comment on these cases and how my work might reshape Anscombe's thesis and the argument from hallucination.

2 Sensory Interaction and Experience

My argument employs the notions of sensory activity and sensory interaction. Neither is exact, but both should be straightforward enough for discussion. As I'm using the terms, a *sensory activity* is any activity of an organism facilitated by its use of a sensory system. A *sensory system* is an apparatus for transducing energy from the environment into signals which can be used by the organism. The nature of sensory systems is complex (see Macpherson 2014), but I'll assume we can take for granted typical examples like the visual system.

Sensory systems trade in information-encoding neural signals (Dayan and Abbott 2001, 8). Millikan's idea (2004) of a producer-consumer distinction is helpful. An organism is a sensory-signal consumer which uses the outputs of sensory systems (which are signal producers) to do, or facilitate, certain things: what I'm calling sensory activities.[2] Due to cross-modal effects (e.g., O'Callaghan 2008) and cognitive penetration (e.g., Macpherson 2012) the

coupling between sensory consumer and sensory system is overlapping, multimodal, and bidirectional, but these are complications that won't matter here.

A paradigm example of a sensory activity is what Matthen (2014) calls *sensory exploration*. Imagine sighting an apple, walking up to it, and picking it up to inspect it. You focus attention on it, look over it carefully, and rotate it with your hand. As you do so you are visually exploring the apple. The act of exploration is the activity you do which is facilitated by your visual system—the sensory activity.

When a sensory activity involves interaction with an object, I'll call it a *sensory interaction*. For example, your visual exploration of the apple involves interaction with that apple in a robust sense: you orient your body movements to it (e.g., walking towards it and turning the apple relative to your body) and physically grasp it in your hand. Perhaps more importantly, what's happening in this interaction is that information about the apple picked up by your sensory systems enables your apple-directed movements. As I'm using the term, not all sensory interactions are quite so robust. For example, if you had merely caught a glimpse of the apple and taken the time to resolve it, that would still count as a sensory interaction. Even if you didn't then use that resolution to guide any particular movements, you still would have used visual information to get the apple in view. But mere causal interaction doesn't seem sufficient. If you had merely thrown the apple a quick glance without looking long enough to resolve it, the sensory activity—namely, glancing around—would not have made any interesting or meaningful contact with the apple. Your glancing would not be an interaction with it.

I take paradigm examples of *objects* to be typical things like individual apples, chairs, rocks, and people. What's key is that objects are *particulars*, and bear features (i.e., properties). You can sensorily interact not only with typical objects like apples, but also the parts of such objects, their surfaces, events, sounds, smells, flavors, ephemera like shadows, and absences like holes (Johnston 2011, 174–76). I will use the term 'object' to include any feature-bearing particular with which you can interact through your sensory systems.

Using sensory organs to *explore*, and using eyes to *glance about*, are two examples of sensory activities. Here are three more: (1) visually identifying locations of distal objects in egocentric space as targets of possible bodily action (Matthen 2005, chapter 13), (2) tracking distal objects (or their features)

through changing proximal stimulation (Burge 2009), and (3) discriminating those objects from the background (Siegel 2006; Schellenberg 2019). So, *locating*, *tracking*, and *discriminating* are sensory activities. They are also (unlike glancing about) plausibly sensory interactions with the objects located, tracked, or discriminated. These activities need not be deliberate; e.g., most people don't intend to track distal objects through changes in proximal stimulation.

Normally there's 'something it is like' for an organism engaged in sensory activity (Nagel 1974, 436). These phenomenal states are typically called 'experiences'. I'll use the term 'sensory experience' specifically for the phenomenal states generated by sensory activities.[3] Although it was once widely held that sensory experiences are always experiences of features (or "sensory qualities") arrayed in space, it's now widely held that experiences are fundamentally organized around objects (Dickie 2010; O'Callaghan 2016; Matthen 2019). What it's like for you, as you use your sensory systems, involves the presentation of particular objects. As you use your sensory systems you seem to find yourself simply confronted with particular feature-bearing objects, objects which seem to intrude into your stream of consciousness (Hellie 2014, 247). These presented objects are *what you experience*, or (alternatively put) what your experience makes you *aware* of. When I speak of the 'objects of experience', I'll always be referring to whatever particular feature-bearing object the experience (seems to) present to its subject. Historically this term has a broader meaning, so that *whatever* an experience makes one aware of is its "object", even if that thing is an object's property or a spatial array of private sensory qualities. While the features (i.e. properties) of presented objects are among the things of which an experience makes us aware, and hence are "objects of experience" in this broader sense, I'll always use the term to refer specifically to the particular feature-bearers presented in experience and around which the experience is organized.

With those preliminaries, we have:

Premise 2: If an experience is generated by a sensory interaction with some object O and seems to be an experience of an object O*, then O* is O.

The idea is that if you introspectively seem to be experiencing some object, and your experience in fact is generated via sensory interaction with a physical object, then what you seem to be—and are—experiencing *is* (numerically identical with) that physical object. Here and throughout I intend the term

'seem' in a phenomenal sense. So, to a rough first approximation, you introspectively seem to be experiencing some object iff what it's like, from your "inner" point-of-view, involves the presentation or appearance of an object. I also intend the term inclusively, so that seeming to experience an object is compatible with actually experiencing it.

The initial motivation for this premise is that the objects of experiences seem to be the very objects with which we interact in the sensory interactions generating those experiences. Consider how things intuitively strike you as you visually explore (i.e., look over) an apple you're holding. Intuitively, it introspectively seems that you are experiencing the apple you're holding and looking at. Your experience strikes you, from a naïve standpoint, as a simple and direct *revealing* or *disclosure* of the apple (Johnston 2006), as if your senses are somehow bringing the apple itself into your private consciousness. This is at least in part due to how your experience is coupled with your interaction. Changes in the appearance of the object of your experience predictably covary with your sensory activity. For example, by turning the apple in your hand you seem to bring a new side of the object of experience into view. If your experience didn't depend on the sensory interaction, then it (presumably) wouldn't be generated by it at all.

As the name suggests, naïve realists about sensory experience take up this intuitive picture (e.g. Campbell 2002; Martin 2004; Fish 2009). They hold that sensory experience is, at the level of metaphysical structure, a relation between the sensing organism and experienced object. The naïve realist need not say that this relation is especially simple or direct (or a "raw" revealing), but the idea is that what it is to sensorily experience a distal stimulus is to stand in some relation to it, e.g. perhaps a relation we'd describe as *perceiving* or *being acquainted with* (Fish 2009, 15; Logue 2012b, 221). Noë (2004), a sensorimotor enactivist who adopts naïve realism, explicitly connects this experiential relation to our sensory interactions. He would say that as we sensorily interact with a distal stimulus, we appreciate (implicitly) how movement affects the way things appear to us, and that this appreciation constitutes experience of (i.e. reveals) the stimulus. Although naïve realists are typically not so explicit, the spirit of their view has it that sensory interaction brings about the experiential relation in such a way that the object you experience is the object with which you sensorily interact. Hence, premise 2 unsurprisingly holds on naïve realism.

Many contemporary philosophers of perception reject this naïve view of sensory experience. Instead of taking experiences to be a revealing of distal

stimuli by the senses, they take them to be representational states. As you, for example, look over an apple you're holding, the neural activity of your visual and tactile systems comes to be in a state which represents the apple. The resulting sensory experience either is this neural representational state, or supervenes on it. Importantly, on this view *what* you experience is what's represented by the experience. While your experience is a representation, your experience isn't what you experience; you, for example, experience the apple itself, not a representation of the apple (Dretske 2003, 73; Crane 2006, 136; Clark 2012, 767; Genone 2016, 6). The standard ways of working out this approach will yield premise 2.

Consider, for example, Dretske's view (2003). Dretske distinguishes between two components of the content of conscious sensory representational states (i.e., sensory experiences): the particular object represented by a given state token, and the properties attributed to that object in virtue of the state token's type. He gives an informational or tracking semantics for the second component, but it's the first that matters here. He says (2003, 68), 'Following fairly standard causal thinking, I take the object(s) of a representation [like a sensory experience] to be the object(s) that stand(s) in the right causal relation to it.' The idea is that a token sensory state (e.g., the sensory experience you now enjoy as you look over your apple) represents whatever particular object caused that tokening in the right way. My point is simple: if an experience is generated by a sensory interaction with an object and seems to present some object to you, then the object with which you interacted will be the one which causes your experience in the right way for the experience to be a representation of it. Hence, what the experience seems to present to you will be, on Dretske's view, the object of your sensory interaction. So, we have premise 2.

Dretske's view isn't idiosyncratic, but instead is closely related to (or perhaps a version of) the popular view that sensory experiences have demonstrative (Burge 2005) or gappy (Schellenberg 2010) content. Demonstrative and gappy contents are context sensitive, so that the particular object represented by a representational state is determined by the situation in which the state (e.g., the sensory experience) is tokened. Assume that a sensory experience is generated by a sensory interaction with an object. That object would be the contextually determined referent of the experience, on these views. Hence we get premise 2.

So premise 2 turns out true on popular representationalist accounts. This shouldn't be surprising. Most representationalists accept the particularity of

sensory experience—that we experience particular objects—and explain that particularity by appeal to representations of those particulars. *Some* explanation will have to be given about how conscious sensory states (i.e., sensory experiences) come to represent the particular objects perceived, and it's hard to imagine any explanation which doesn't somehow ground that particular content in sensory interactions. Matthen, for example, explicitly grounds the particular content of visual experience in the sensory activity of locating objects in egocentric space (see Matthen 2005, 300–305). Even representationalists who deny that sensory experience has particular content (e.g., Tye 1995; Pautz 2009) usually accept the particularity of the experience; they simply explain that particularity in some other way—usually by appealing to causation (e.g., Pautz 2009, 499, 2010, 286). But a causal explanation of the particularity of experience will also bring us back to premise 2, since sensory interaction necessarily involves the distal object causing the experience.[4]

A potential problem for premise 2 is that a given sensory experience is usually generated by a variety of different sensory interactions. This can take many forms. Consider cross-modal illusions like the McGurk effect (Tiippana 2014) and the sound-induced flash illusion (Shams et al. 2000). For instance, in the sound-induced flash illusion your visual experience of flashing disks is affected by hearing beeps. So we have a counterexample to premise 2: your visual experience of the disks is generated (in part) by your auditory interaction with the beeps, but the visual experience isn't an experience of beeps.

This sort of counterexample doesn't reveal any fundamental problem with premise 2. It does not show that sensory interactions are disconnected from object-presenting sensory experiences. What it shows is that the object of a sensory experience will be *one* of the objects of the generating sensory interactions. What's needed to fix premise 2 is a *cateris paribus* clause which excludes these sorts of cases. The new premise 2 would be something like: If an experience is generated by a sensory interaction with some object O, seems to be an experience of an object O*, *and there isn't some other object with which the organism interacts sensorily which is O**, then O* is O. For example, in the sound-induced flash illusion you have an auditory interaction with beeps, but these are not the objects of your visual experience, since your visual experience is of the flashes.

When amended in this way premise 2 doesn't carry much explanatory power, but is true. The lack of explanatory power won't matter in the ar-

gument to follow, because in the relevant cases of hallucinations there won't be potential objects of experience which aren't in the past. The thesis being defended here is that some hallucinations are experience of past-perceived objects, and premise 2 will suffice for this purpose so long as all the candidate objects of sensory interactions generating hallucinations are all located in the past. Hallucinations are generated by sensory activities which fail to interact (in the normal way) with any objects in the present here and now; that's what makes them hallucinations. So if hallucinations are generated by sensory activities which involve interaction with distal physical objects, those objects must either be in the past or (if in the present) must be causally connected to the subject in some deviant way. None of the three cases I'll discuss involves deviant causal interactions with concurrent objects of experience.

3 Plasticity, Pattern Completion, and Recall

Premise 2 is fairly straightforward. The more controversial part is premise 1: the claim that hallucination involves sensory interaction with past-perceived objects. The argument for premise 1 starts by giving two jointly sufficient conditions for sensory interaction:

Premise 3: A sensory activity of an organism is an interaction with an object if (1) that object is a causal antecedent of the neural signals mediating the activity, and (2) the use or function to which the organism puts those signals involves that object in the right way.

The notion of a sensory interaction is not rigorous. At this point in the discussion we are conceptually exploring its bounds given paradigm examples of it (like sensory exploration) and its role within theorizing (e.g., how it grounds reference to particulars within representational accounts). Premise 3 is a proposal for another step in filling out the notion of sensory interaction. We should understand sensory interaction in a way which makes the two proposed conditions jointly sufficient for it. There are three things which can be said in favor of this proposal.

First, premise 3 is an analysis, into two separable conditions, of the basic idea that you interact sensorily with an object when you use information from or about that object (picked up through your sensory systems) for some appropriate action directed at it. This basic idea came out above when it was noted that, when you interact sensorily with an object, you use information

from or about that object picked up through your sensory systems. Any sensory activity which satisfies the two conditions in premise 3 will plausibly be such a case, i.e. will plausibly be a case in which the organism doing the sensory activity uses sensory-derived information from an object in a way that somehow involves that object.

Second, paradigm examples of sensory interaction satisfy both conditions given in premise 3. As you visually explore the apple in the example given above, the apple is a causal antecedent of the neural signals facilitating your exploration and you put those signals to use to guide your approach to the apple, to enable your grasp of it, and to assess its various qualities like ripeness and color. In sensory exploration you are concerned with and act towards an object in a way that's facilitated by neural (information-carrying) signals caused by that object.

Third, if a sensory activity meets the two conditions in premise 3 and generates a sensory experience, then the resulting experience is an experience of the relevant causal antecedent. That is, the considerations in favor of premise 2 given above still go through if we assume premise 3. The argument that premise 2 fits with common versions of representationalism and naïve realism is itself only driven by the two features of sensory interaction mentioned in the two conditions of premise 3: that it involves causal and functional connections.

Let's consider premise 4.

Premise 4: At least some hallucinations (and probably all normal perceptual experiences) are generated (in part) by sensory activities which are mediated by neural signals which have past-perceived objects as causal antecedents.

The basic considerations in favor of this premise run as follows. The neural activity mediating normal sensory activities—whether resulting in hallucinatory or successful perceptual experience—is prompted by changes in the membrane potential of sensory receptors due to stimulus energy. But the exact shape of this mediating neural activity (e.g., spike-rate frequency and the spatial pattern of spike-rate frequency changes across patches of cortex) is determined not just by the stimulus-driven changes in receptor membrane potential, but also by the weights of the synapses connecting the neurons in the relevant circuits. These synapses exhibit *plasticity*. Their weights are (in turn) shaped by past activity of the system.[5] In the normal case, at

least some of that past activity was (of course) driven by past sensory interactions, interactions with past-perceived objects. So past-perceived objects causally affect the neural activity mediating current sensory activities via the past neural activity they prompted and effects that past neural activity had on synaptic weights. This reasoning won't go through for at least one case: the case of a fully envatted brain which has always and only been fed input based on artificial computer simulations. But it will go through for the hallucinations of a subject who usually is in normal perceptual contact with the physical world.

Given the causal structure just outlined, it's hard to deny premise 4. The real work, I anticipate, is in premise 5. Even common sensory interactions, like sensory exploration, leading to successful perceptual experiences have the sort of past-pointing causal connections highlighted by premise 4. But these are not interactions with those past causal antecedents; they are interactions with the concurrent causal antecedents affecting sensory receptors: e.g. the present distal objects being explored. Unless we're to beg the question against the possibility of having sensory interaction with the past, this is presumably to be explained by the *use* to which the organism is putting the mediating neural signals in the sensory activity. Sensory exploration (for example) is an interaction with the present object (and not any past causal antecedent) because the purpose of the exploration is present-focused; the organism is sensitive to the present object, orients its exploratory behavior with respect to it, and that activity functions to gather information about what's here now.

So if the sensory activities which generate hallucinations are interactions with past causal antecedents, it is because the generating activities are past-focused. This brings us to premise 5:

Premise 5: In at least some of these hallucinations (generated by sensory activities mediated by neural signals with past-perceived objects as causal antecedents), the generating sensory activities involve those past-perceived objects in the right way.

What it is for a sensory activity to 'involve' some bit of the world 'in the right way' is imprecise, but the above discussion should already indicate some of what's meant. There's at least one sensory activity which clearly involves past-perceived objects in the right sort of way: recall.

Perhaps the best example of recall is episodic memory (see Tulving 1983; 2002). In episodic memory one not only recalls some bit of information—that

such-and-such happened or is true—but "relives" or re-experiences a previous sensory experience via mental imagery. For example, I can call to mind select moments from breakfast this morning. This recall is *past-pointing*. I recall a specific past instance of preparing food and eating it. This recall is an activity which engages with particular past-perceived objects: e.g., the eggs I cracked into the pan and the cheese I spread on my bagel. I recall them not only in tenseless abstraction, but as they were in that past event. I am recalling certain past temporal slices of those things. The recall I'm doing right now as I think about this episodic memory is a (sensory) activity which involves these past-perceived objects in the right sort of way for interaction.

My suggestion is that in at least some cases of hallucinations (generated by sensory activities mediated by neural signals with past-perceived objects as causal antecedents), the generating sensory activities involve recall. Since recall is past-pointing in the right sort of way, premise 5 will follow immediately. My argument for this suggestion will turn on the following premise:

Recall Premise: If a given sensory activity is facilitated in part by neural pattern completion, then that activity involves a form of recall.

The standard account of episodic memory provides an introduction to neural pattern completion.

The net effect of synaptic tuning varies (see Feldman 2012), but often the result is that a pattern of activity primes a neural circuit to repeat that pattern when provided part of it as input (Jackson 2013). The standard view is that this pattern completion is what facilitates episodic memory: episodic memories are recalled by repeating the patterns of neural activity which facilitated the original (recalled) sensory interactions (Brogaard and Gatzia 2017, 9). The idea is that the original pattern of activity tunes the synaptic weights in relevant memory-storing areas, leaving behind a memory "trace" or "engram" in the circuit (Liu et al. 2014, 59). These memory traces are the effects of the original pattern—presumably the tuned synaptic weights—which allow the organism to repeat that pattern by inputting a part of it.

If pattern completion facilitates recall in the case of episodic memory, perhaps it's facilitating recall in other cases as well—hence, the recall premise. The next step is to argue that the sensory activity generating some cases of hallucination is facilitated by pattern completion. Aside from dreams, philosophers often have in mind speculative examples of hallucinations, like the life-long fully envatted brain mentioned above or a case in which the would-be perceived object is removed while a subject's neural activity is

maintained by direct stimulation (e.g., Valberg 1992, 9–11; Martin 2004, 40; Logue 2014, 198). Neither of these are *prima facie* cases of pattern completion. But hallucinations generated by misleading visual cues provide a good case.

A well-known example is seeing subjective contours, as in Kanizsa's triangles (see Kanizsa 1976). In these cases while viewing an image you visually experience contours, or edges, of some shape which isn't really there. You hallucinate not only the edges of the shape, but the shape itself. That is, as you look at the image you seem to experience not just the contours or edges of a triangle, but a triangle.[6] A Kanizsa triangle is an hallucinated object (feature-bearing particular), and hence these hallucinatory experiences are relevant to the main argument of this paper.[7] It should be immediately plausible that some sort of pattern competition is behind these hallucinatory experiences. While the details aren't critical (so long as some sort of pattern completion is involved the recall premise will go through for this case and we'll have completed the argument) it's worth examining one way the details could be unpacked.

For this I want to turn to Albright's account (2012) of visual processing. In contrast to a purely feed-forward extraction approach, Albright (2012) suggests (in line with predictive processing approaches) that overall activity in visual processing areas results from a mix of bottom-up and top-down input. This top-down input originates (on Albright's view) in the inferior temporal (IT) cortex, which he suggests serves in part as a 'long-term repository of visual memories', with the medial temporal lobe and hippocampus serving a role in the formation of these memories (2012, 229). Crucially, the top-down input from IT cortex seems to be a form of pattern completion similar to that facilitating episodic memory recall: repeated exposure to concurrent stimuli strengthens connections in IT cortex so that exposure to one stimulus prompts top-down signals which trigger activity normally associated with the missing stimulus. On Albright's view, the neural circuits facilitating this top-down generation of visual activity are the same ones involved in mental imagery tasks. The extent to which bottom-up or top-down input dominates overall visual activity depends on the quality of the bottom-up signal; when input from sensory receptors is incomplete, noisy, ambiguous, or otherwise defective, input from this top-down imagery system kicks in to fill in the gaps (see also Tang et al. 2018).[8] How we fill in missing information during perception isn't different from what we do when we generate mental images of past-viewed objects based on episodic memory traces (Albright

2012, 237–39).

The kinds of misleading visual cues which prompt hallucinations (as when viewing Kanizsa triangles) are certainly examples of ambiguous or incomplete input which the visual system might fill-in via pattern completion. The visual cues now prompting your hallucination of a triangle were in the past associated with triangles. That association prompted activity which left visual memory traces (perhaps in the IT cortex) which are reactivated when viewing the cues.

Albright provides a second example of a hallucination generated by top-down completion:

> Consider, for example, the "vanishing ball illusion": in this simple yet compelling trick, the magician repeatedly tosses a ball into the air. On the final toss, the ball vanishes in mid flight (...). In reality, the ball never leaves the hand. The illusion is effected by the use of learned cues that are visible to the observer, including the magician's hand and arm movements previously associated with a ball toss, and the magician's gaze directed along the usual path of the ball. The observer's inferences about environmental properties and events are probabilistically determined (from the associated cues) but the inferences are incorrect. According to the implicit imagery hypothesis, these flawed inferences are nonetheless manifested as imagery of motion along the expected path. Moreover, this imaginal contribution to perceptual experience is likely to be mediated by top-down activation of directionally selective MT neurons, in a manner analogous to the effects [observed in Rhesus monkeys]. (Albright 2012, 238, reference omitted.)

Albright is at the end referring to an experiment he completed (see Schlack and Albright 2007) using Rhesus monkeys which shows that motion-sensitive neurons in visual cortical area MT can be trained via association to respond to (artificial) cues for motion, such as arrows (Albright 2012, 230–31). As Albright describes the results, 'evidence indicates that the learning-dependent responses to arrows in area MT are ... a cued top-down reproduction of the activity pattern that would be elicited in MT by a moving stimulus projected upon the retina' (Albright 2012, 232). In the case of the vanishing ball illusion, on the illusory toss the observer has a visual experience of a ball leaving the magician's hand, but there is no such ball.[9] This hallucinatory

experience (of a ball which isn't there) is generated (at least in part) by the cued reproduction of visual activity that was previously elicited by the prior real tosses.

So Kanizsa triangles and the vanishing ball illusion plausibly provide two cases of hallucinations generated by sensory activities facilitated by pattern completion. It might be objected that the recall premise is false: pattern completion need not be facilitating recall. For example, one proposal is that the neural circuits which facilitate episodic memory also enable prospection (planning for and imagining the future), imagining counterfactuals, and taking others' viewpoints (Addis et al. 2007; Buckner and Carroll 2007; De Brigard 2014; Michaelian 2016); in these cases, pattern completion presumably facilitates what Buckner and Carroll call *simulation*, or the 'construction of an imagined alternative perspective' (2007, 49). Returning to Kanizsa triangles and the vanishing ball illusion, at first glance it doesn't seem like recall is involved either. In response, let's start with the question: what is the sensory activity in which you're engaged as you view the illusory toss or view the image prompting the experience of a Kanizsa triangle?

In the first case you are visually tracking the ball's position. In the second case you are visually exploring the medium of the image (a screen or piece of paper) and the marks forming the image. Both activities plausibly involve doing other things as well, e.g. discriminating the ball and marks from the background and tracking their shape through changes in proximal stimulation. When viewing the image this discrimination and tracking plausibly happens for the sake of something else: *sorting*. As Matthen suggests (2005, 13–35), organisms use sensory systems to sort distal stimuli into similarity spaces along dimensions which are relevant to them. Some shape is depicted in the image, and as you view it you use your visual system to classify the depicted shape.

Whatever the sensory activities involved in these cases, my opponent will press that the underlying pattern completion is facilitating present-focused activities like sorting and position-tracking, not recall. I agree that sorting and position-tracking are what's being done in these cases, but this is compatible with the subject also doing recall. When faced with incomplete or ambiguous sensory-receptor input (as in these cases), you classify the shape of the distal object or track its position by first recalling similar past-viewed objects or similar past trajectories. When given defective stimulus input, sorting and position-tracking proceed by first doing recall. The reason we should think that recall is first done to fill in the gaps is because of the

character of the facilitating neural activity: it involves mnemonic pattern completion.

Something very much like this suggestion is hinted at by Buckner and Carroll in their own discussion of prospection and imagining alternative viewpoints. Although they say that what the relevant neural circuits are doing in these cases is simulation, it's simulation 'based on our past experiences' which relies 'on autobiographical information' (2007, 49). They suggest (2007, 52) that 'Simulations of others' perspectives, and of ourselves in another time, might be built on specific past instances, as captured through medial temporal processing.' Responding to the same findings regarding the overlap of episodic memory and prospection, De Brigard similarly suggests that the cognitive system behind episodic memory actually functions to recombine 'encoded traces into representations of possible past events that might or might not have occurred, presumably in the service of constructing mental simulations of possible future events' (2014, 158). So the point that pattern completion facilitates prospection and other non-past oriented perspective shifts is not a challenge to the recall premise, as it might allow for these activities by first enabling recall (of past-perceived objects that are recombined into the prospected scene).

It might be asked if the above considerations undercut the idea that hallucinations are *sensory* experiences. Don't these considerations show that hallucinations are a form of imagination?[10] I think this is exactly right: the considerations raised in the discussion of the recall premise strongly suggest that hallucinations (at least, the ones under consideration here) are a state of imagination. But this outcome isn't in tension with my main thesis. The central claim of this paper is that some hallucinations (e.g., those prompted by misleading cues) involve *experience* or *phenomenal awareness of past-perceived objects*. The emphasis of the thesis is on how these hallucinations involve awareness of the past. The nature of these experiences, e.g. whether they are a form of imagination, is orthogonal to this emphasis. I have said that these are *sensory* experiences, but by 'sensory experience' I just mean any phenomenal state generated by a sensory activity. Since the neural machinery of our sensory systems is shared between imagination and perception, it's not unexpected that some states (e.g., hallucinations prompted by misleading cues) could result from activities of our sensory systems while still being a form of imagination.[11]

4 Two More Cases of Past-Pointing Hallucinations

Two other cases of hallucination are plausibly generated by sensory activities involving recall, and so satisfy premise 5 and provide examples of experience of the past. The first is Charles Bonnet Syndrome. CBS is a condition in which those with vision loss due to retinal or optic-nerve damage suffer hallucinatory visual experiences in the period immediately after the onset of the loss. As ffytche describes it (2013, 50), 'The hallucinations occur as a series of discrete episodes lasting seconds or minutes with, initially, many such episodes every day but a graded reduction in their frequency over time.' These experiences often include people and everyday objects, but also disembodied faces and simple shapes or patterns (ffytche 2013, 51). The hallucinations are often bizarre, e.g. including oversized objects or faces with distorted features. For those with partial vision loss, the hallucinated objects are typically experienced as situated within the actual scene seen in front of them; the hallucinated objects are experienced vividly, even if the actual objects perceived are experienced as blurry (ffytche 2013, 50).

Albright (2012, 239) suggests that CBS hallucinations result from the same top-down image generation which completes defective input in cases of misleading-cues. Assuming the recall premise, CBS hallucinations are generated by recall. But the cause of CBS hallucinations is a point of contention in the literature. On the basis of his earlier, seminal fMRI work (ffytche et al. 1998), ffytche argues (2013, 53–4) that CBS hallucinations do not originate from the same top-down neural circuits which facilitate mental image generation. This work showed that CBS hallucinations correlate with activity in the fusiform gyrus, and in ways that fit well with the patterns of activity usually prompted in healthy patients when viewing similar stimuli. For example, colored CBS hallucinations correlate with activity in areas typically associated with color processing and hallucinations of faces correlate with activity in areas typically active when viewing faces. In contrast, CBS hallucinations were not found in his work to be associated with areas he says are related to mental imagery generation, specifically the frontal, parietal, and medial temporal lobes (ffytche 2013, 53).

Instead, ffytche favors a deafferentation hyperexcitability model. CBS hallucinations often arise after macular degeneration (damage to the macula, a part of the retina including the fovea), which causes a loss of input from

the fovea and adjacent areas but leaves intact the periphery. This damage cuts off input to visual neurons with receptive fields corresponding to the damaged area. Without input, it's thought that these deafferented neurons' receptive fields enlarge over time until they overlap with the periphery enough that peripheral stimuli lead to activity in them (see Painter et al. 2018, 3475). A recent EEG study by Painter et al. (2018) has found evidence that macular degeneration patients who suffer CBS hallucinations do exhibit deafferentation hyperexcitability in visual processing areas, especially the early visual areas V1 and V2 (2018, 3477), while this hyperexcitability is not found in macular degeneration patients who do not suffer CBS hallucinations.

All sides agree that the neurobiology of CBS hallucinations is poorly understood, so at best the above is preliminary data for the debate. There are three points to make in Albright's favor. First, Painter's study observed deafferented hyperexcitability without accompanying CBS hallucinations. So although the subjects who exhibited the deafferented hyperexcitability were those (and only those) who also suffer CBS hallucinations, peripheral stimulation which sufficed to induce this hyperexcitability was not sufficient to induce hallucinations. If hyperexcitability was the main mechanism behind the hallucinations, it's unclear why it would be observed without them.

Second, deafferented hyperexcitability was most pronounced in early visual areas, while CBS hallucinations are associated with activity throughout the fusiform gyrus. Painter et al. used a simple checkerboard pattern, so it's possible that a more complex stimulus (e.g., a face) in the periphery would have prompted stronger activity outside early visual areas. But in general CBS hallucinations are not closely tied to the presentation of similar stimuli.[12] In fact, often CBS hallucinations are prompted by loose, abstract associations: e.g., reporting the work of others, Albright notes that 'upon hearing an account of the revolutionary war, one patient with CBS reported a vivid percept of a winking sailor' (2012, 239). So the pattern of deafferented hyperexcitability does not fit CBS hallucinations, and those hallucinations aren't tightly tied to peripheral stimuli anyway.

Third—and most important—ffytche assumes that a lack of associated activity in post-visual memory-related areas (such as the medial temporal lobes) means that CBS hallucinations are not generated by top-down imagery circuits; but this assumption is open to serious doubt. Albright's view is that these post-visual areas are important for memory formation, but that visual memories are themselves stored (via pair coding) in the IT cortex (Albright 2012, 230). As van de Ven and Sack (2013) note, this view is supported

both by classic experiments showing that hallucinations based on episodic memories can be prompted by stimulating sensory cortex (e.g., Penfield and Perot 1963) and by more recent work with transcranial magnetic stimulation (TMS). The lack of activity in post-visual circuits might actually fit well with the nature of hallucinations; one suggestion is that these post-visual areas facilitate perspective shifting (see Buckner and Carroll 2007), and a hallmark of hallucinations is that they present themselves as experience of the here and now (and not as simulations which are temporally or spatially shifted).

So ffytche does not have a compelling case against the view that CBS hallucinations result from mnemonic pattern completion. It's plausible that CBS hallucinations are prompted by the activation of visual memory traces (whether or not in the form of pair coding) stored in visual processing areas. These traces are perhaps activated by residual input from the visual periphery that's distorted by deafferented hyperexcitability. Just as feedback from visual memory traces prompt the hallucinations of Kanizsa triangles and the disappearing ball by completing patterns cued from those incomplete inputs, residual defective input in CBS patients prompts pattern completion from visual memories.

What is the sensory activity generating CBS hallucinations? Presumably it's just the same sort of activity generating normal visual experiences: e.g., tracking, discriminating, or sorting. But like in the misleading cues cases, the receptor input is defective in some way which prompts mnemonic pattern completion; in effect, the subject attempts to track, discriminate, sort, etc, despite defective input by first doing recall.

Dreams are the final potential example of hallucinations generated by sensory activity which involves recall. It is perhaps obvious that dreams often in some way involve, or depend on, memories. For example, often in dreams you find yourself in a familiar place, or seeing familiar people. This familiar imagery is often more specific: it is what Freud called 'day-residue', or 'elements that [connect] with experiences of the previous day' (Nielsen and Stenstrom 2005, 1286). As Nielsen and Stenstrom note in their review of dreams and memory (2005, 1286), 'When psychophysiological methods were introduced into the study of dreaming, ... Results confirmed the robustness of day-residue memory elements and largely attributed these elements to the fragmentation and transformation of episodic memories.' So the idea is that the hallucinatory experiences composing dreams are often made up of rearranged parts of past perceptual experiences, via the fragmentation and

transformation of episodic memories (see also Windt 2015, 546–49). For example, in one study 65% of dream reports were found to include day-residue, with 1-2% of reports including complete episodic memories (Fosse et al. 2003; cited in Windt 2015, 604).[13]

It is tempting to infer that dreams involving repackaged episodic memories are generated by recall. If the involvement of pattern completion can be inferred from the presence of memories in dreams, the recall premise might be used as well. As before, an opponent might object that just because a dream is generated by the reactivation of an episodic memory trace (or by the binding of many fragments of such traces) does not mean that what the organism is trying to do by reactivating that trace is recall something. The fact that we typically don't experience memory-based dreams as mnemonic, but instead as happening now (Windt 2015, 547), is one reason to resist. A second problem, which could have been raised for CBS hallucinations, is that if the aim were to recall, why mix and match fragments of memory traces?

But, as we've already seen, dreams might be generated in part by recall without that recall aiming at complete and accurate episodic memories. Misleading-cue and CBS hallucinations are generated by typical sensory activities (e.g., tracking and sorting) that are facilitated by recall. My suggestion is that whatever the generating activity in the case of dreams, that activity is likewise facilitated by first doing recall. What is the generating activity?[14]

A natural suggestion is that dreams result from the same sort of activity involved in prospection and other perspective-shifting: simulation. For example, Windt's empirically well-supported view is that as you dream, what's happening is that you're simulating an immersive spatiotemporal world centered on yourself (Windt 2010; Windt 2015, 568). Windt's full view follows the empirical work of Blumberg (Blumberg 2010, 2015; Blumberg et al. 2013) and the predictive processing framework of Clark (2012) and Hobson (Hobson et al. 2014; Hobson and Friston 2012, 2014). In a nutshell, the idea is that during sleep the brain engages in what Windt calls *bodily self-sampling* (Windt 2018, 2603). This involves letting selective motor commands through the motor output blockade in the brainstem while trying to predict the interoceptive bodily feedback generated by the resulting muscle twitches.[15] Error from these predictions allows the brain to improve its body maps for better sensorimotor coordination while awake. Along with predicting the interoceptive signals—modeling bodily action, essentially—the brain creates models of an external environment which fit that bodily action (Windt 2018,

2616–18). Bodily simulation generates tactile and movement sensations in dreams, while simulation of an environment generates experience in exteroceptive modalities like vision.[16] Models of an external environment are created via some mix of episodic memory, imagination, desire, and (nonsensory) expectation (Windt 2018, 2617).[17]

We can accept something like Windt's account of dreams being generated by simulation without rejecting that recall is involved. Dreams aren't generated by pure episodic memory recall, but the fact that they seem to be stitched together from bits of episodic memories suggests that we are simulating an environment by first recalling past-perceived objects to populate that environment. More speculatively, presumably the neural activity facilitating this simulation starts as partial signals (from the bodily self-sampling, unblocked exteroceptive receptor input, or cognitive processing areas) which prompt pattern completion from stored sensory memory traces.[18]

5 Objections and Replies

This section collects together some remaining objections and replies.

Objection 1: I'm claiming, or assuming, that *recalling an object* is a sensory interaction with it, whether that recall is part of tokening an episodic memory or generating a hallucinatory experience. An anonymous referee, echoing several other commentators, points out that 'The effect of the (past) world on the perceiver is entirely one way, unlike a standard sensory relation to the world.' The worry is that the one-way nature of recall means that it's not a type of *interaction*.

Reply: While there's a fair question about whether the term 'interaction' is really the best to describe the metaphysics, what's crucial for my argument is whether the connection between recalling organism and recalled past-perceived object is of the sort satisfying premise 2. What makes standard sensory interactions, like sensory exploration, relevant to the identity of an object presented in experience isn't (so I think is plausible) the concurrent corporeal exchanges between them, *but the flow of information*. As I emphasized in discussing sensory interactions, what makes for a sensory interaction (of the relevant sort) is that the organism uses information from the distal object to facilitate some object-directed action. Recall satisfies this condition (whether it's recall done in token episodic memories, or recall done in hallucination). In recall the organism uses information from the distal object

to facilitate some object-directed action, whether that action be to imagine the object, judge some of its previously observed properties, communicate about it to others, fill in gaps in current sensory input, fill out a simulation, or something else. The upshot is that recall has the features of more standard sensory interactions relevant to premise 2, i.e. relevant to identifying the objects of experience, and hence counts as a sensory interaction for my purposes here.

Objection 2: My emphasis on sensory activities (in the sense outlined in the section 2) is misplaced and the recall premise is beside the point. Take Kanizsa triangles. A typical computational explanation—likely on many reader's minds—would say that shape cues in the marks on the image prompt your visual system to extract or construct the shape of a triangle (the resulting neural triangle representation leading to your hallucinatory experience). This explanation tightly ties sensory experience and its content to representational states of sensory systems, and tightly grounds that sensory state content in neural representations (e.g., Tye 1995; Prinz 2000, 2006, 2011). For this theorist, what's really important are the neural representations computed by the sensory systems. While pattern completion might be the underlying neural mechanism realizing these representation constructions, the repeated pattern isn't encoding (in this case) memories, but instead a model or prediction.

Reply: I've framed things in a way which leaves open both this computational approach and approaches on which experience supervenes on the whole sensory interaction (such as Noë's sensorimotor enactivism). If this view is right, then my talk of a sensory activity generating experience should simply be reinterpreted as talk of the mediating neural activity generating the experience. More importantly, the overall sensory activity itself matters on either view. Something must fix the content of neural representations, and it's generally accepted that a pure causal or informational theory won't work. If some bit of neural activity represents a triangle, it's not merely because it was caused (or is usually caused by) a triangle, or that it carries information about triangles. As Dretske (1988) and Millikan (1984) note, it must be something about the use, or function, to which those signals are put which makes them representations of triangles.

The upshot is that there must be something which makes the neural activity resulting from viewing the image a representation of a triangle—and that will have something to do with both the causal antecedents of that neural activity and the organism's overall sensory activity. If my proposal for the

sensory activity is correct, the resulting neural triangle representation which leads to the hallucinatory experience will be a representation of some past-viewed triangle, and hence the hallucinated triangle will be a past-viewed triangle. My opponent cannot simply assume that the relevant representation constructions are encoding models or predictions which are detached from memory content.

Objection 3: CBS hallucinations typically involve objects not obviously from their subject's past, or atypical objects like distorted faces, and so aren't experience of the past. *Reply:* Just as dream experiences involve awareness of scenes which are stitched together from various past-perceived objects, a novel object in a dream or CBS hallucination might be stitched together from various past-perceived object parts. So, for example, experience of a strange man in a dream or a distorted face in a CBS hallucination might be constituted by experience of various past-perceived arms, torsos, noses, eyes, etc. Hume famously suggested (1993, 11) that all "ideas" are just recombinations of sense impressions. Similarly, my suggestion is that all hallucinated objects in misleading-cue hallucinations, CBS hallucinations, and dreams are just recombinations of past-perceived objects.

Objection 4: Let's say the Humean recombination response works, so that the hallucinated objects in these cases are stitched together from parts of past-perceived objects. Still, the overall hallucinated object itself (e.g., the strange man) won't be a past-perceived object. *Reply 1:* The first possible response is to concede the point: the objects of hallucination in these cases aren't necessarily themselves past-perceived objects, but are merely composed of past-perceived object parts—perhaps in some strong sense of metaphysical or mereological constitution. If these mereological sums of past-perceived object parts exist, then although these cases of hallucination won't be experience of past-perceived objects, they will still be experience of *something*. They won't be experience of things that don't exist.

Reply 2: A second reply is to press back against the assumption that the relevant cases actually involve awareness in the right sense. Say I (seem to) see a strange man in a dream. Is it really the case that I have sensory awareness of some particular object, a strange man? Not all experiences involve awareness of material objects. There's a pen on the desk in front of me. If I throw it a quick momentary glance, or rest my gaze so that it falls in my visual periphery, I'm aware of some of its features (i.e., properties), but not aware of the pen itself (or so I would suggest). I see some colors in certain shapes, but don't make contact with the pen itself or see it as an object.

(For that I need to look more carefully.) It's possible that hallucinatory experience of novel "objects" in these cases is like this experience: when I dream a strange man, I don't have visual awareness of (or even "as of") a strange man; instead, I have visual awareness of some feature-instance collection that has the look of a man.

Objection 5: Is it really the case that all objects of hallucination in these cases can be explained as recombinations of past-perceived object parts? Aren't there some truly novel dream experiences? This objection is raised by Rosen (2018) in a discussion of dreams and Noë's sensorimotor enactivism. As Rosen presses, if, as the enactivist says, sensory experience emerges out of sensorimotor exploration, how could dreams involve sensory experience? Responding to this problem, Noë (2007, 472) has suggested that dream experience is explained by past perceptual experience. He doesn't fill this out with any detail. Rosen suggests a few possible ways to understand Noë's suggestion, including that dreams are generated via memory (Rosen 2018, 308–309). Rosen is also not clear on how the details would go, but she considers the response that dreams 'are most likely simply amalgamations of past experiences of the world' (Rosen 2018, 299) and objects (2018, 309) that dreams 'may involve, for example, bodily sensation or sounds that are unlike anything experienced while awake, or perhaps at a greater intensity, such as intense colour saturation'.

Reply: To start, note that some of Rosen's examples (novel bodily sensations and unrealistic colour saturation) aren't objects, and so aren't immediately germane to the discussion. In a similar way, one might have wondered about Hume's missing shade of blue (Hume 1993, 12), but again that isn't experience of an object. What's needed (to pose a problem for my thesis) is a dream experience of a truly novel *object*, i.e. an object that's not merely an amalgamation of past-perceived object parts. The possibility of dreaming (or otherwise hallucinating) an object *with never before experienced properties* isn't obviously problematic for my view. For example, nowhere have I committed to the claim that experiences are raw, simple, or direct revealings of objects of a sort which could not allow for the misattribution of properties.[19]

But let's allow, for the sake of discussion, that there could be dream experiences of novel objects. Since my argument doesn't assume Noë's sensorimotor enactivism, it's open to me to simply accept that not all dream experiences are experiences of the past. (Premise 2 merely places a sufficient condition on sensory experience, not a necessary one.) My claim has been that *some* interesting cases of hallucination are experience of the past, not

that all are. If there are truly novel dream experiences, then premise 1 will fail for these cases. But that still leaves open that premise 1 is true for lots of other instances of dreaming (and other instances of misleading-cue and CBS hallucinations). I think there are ways for the enactivist (or anyone who wants to hold that sensory interactions are necessary for sensory experiences) to press back against potential cases of truly novel experiences, but that would take us too far afield.

Objection 6: A life-long fully envatted brain will presumably dream as well, and it's possible that it might have dream experiences with exactly the same phenomenal character as your own dream experiences. But this envatted brain has never had any sensory interactions at all, and so these dreams can't be experiences of past-perceived objects. Assuming sameness of phenomenal character implies sameness of the objects of experience, your own dreams aren't experience of the past.

Reply: This objection shows too much; it's just the argument for hallucination applied to this specific case. In any case, anyone who denies the assumption that sensory phenomenology is object-dependent will have an easy time answering this objection. They will simply say that when you or I have the experience at issue, say in a dream, it will be an experience of a past-perceived object, but when the envatted brain has the experience, it will be an experience of something that doesn't exist. This will include most representationalists, who accept that two instances of the same phenomenal character can make you aware of different objects (e.g., Soteriou 2000, 186; Chalmers 2006, 108; Siegel 2010, 169; Schellenberg 2011, 738–740). These representationalists want to distinguish between the total content of a token experience and the part of its content which affects its phenomenal character. They deny that this latter phenomenology-affecting content involves the particular object interacted with (see Speaks 2009, 557; Tye 2009, 258).

Someone who accepts the object-dependence of sensory phenomenology will need a different strategy. This includes (at least some) naïve realists. I see no other option for them but to deny the assumption that a life-long fully envatted brain could have token dream experiences with the same phenomenal character as some token dream experience of a normally situated subject. I think this is a strategy the naïve realist can plausibly press, but (again) there is no room for discussion here. Instead, I will simply note that my proposal in this paper raises no new problem for most naïve realists, who already have to deny that the phenomenology of token experiences generated by sensory interaction with the world can be completely reproduced in token

experiences which lack that interaction. My work makes things easier for this naïve realist, who now can account for many interesting cases of hallucination in the same way they account for perception: as the revealing of distal physical objects through sensory interaction.

6 The Argument from Hallucination

Reflection on hallucination suggests that what we experience isn't what we're interacting with through our sensory systems. Hallucination involves experience of something that's not really there. So it's inferred that in this case you're not experiencing the physical object you seem to be. Hume would have said that instead you're experiencing 'impressions' (1993, 10), Brentano 'intentional inexistents' (1874, 92), Russell 'sense-data' (1912, 12), and many contemporary psychologists 'percepts' (e.g., Crick and Koch 1998, 97). Impressions, intentional inexistents, sense-data, and percepts are all supposed to be nonphysical or merely in the head. Using a transfer principle (e.g., the premise that any perceptual experience could have been hallucinatory or that good and bad cases of sensory experience share the same metaphysical structure), it's then inferred that what we're aware of in even good cases of successful perception are also mind-dependent sense-data (e.g., Ayer 1956, 90–113; see also Valberg 1992; Smith 2002).[20] The result is indirect realism: the view that what we directly experience are mind-dependent objects, with the distal physical world at best indirectly perceived through these.

What I've just summarized is the argument from hallucination. As Crane explains (2006), contemporary philosophers respond to this argument, in an effort to avoid indirect realism, in one of two ways. Following Anscombe (1965), representationalists explain how we could experience something not there by assimilating experience as a kind of representation. Following Austin (1962), disjunctivists (including many naïve realists) reject the transfer step (e.g., Hinton 1967; Snowdon 1980; McDowell 1986). Both Anscombe and Austin framed their work as direct responses to the argument from hallucination and as a way to avoid sense-data theory (a version of indirect realism).

An intriguing possibility raised by my work is that both sides have missed the real problem with the argument. Perhaps the argument goes wrong by misinterpreting what's implied by the nature of hallucination. In short, hallucination can be experience of what's not there while still being an experience of some sensed distal physical object. If, as I've argued, hallucination is ex-

perience of the past, then although hallucinations are experiences of what's not really there, their objects are indeed physical objects with which we interact through our sensory systems—physical objects located in the past. It is usually taken for granted that hallucination cannot be a relation, i.e. that it cannot put us in sensory contact with the world (e.g., Logue 2012a, 173; Hellie 2014, 255; Moran 2019, 10), but my work here suggests that this is too quick.

Of course, I haven't argued that *all* cases of hallucination are experience of the past and my argument won't extend to the cases we've left out without significant work. But if it's convincing that some interesting cases of hallucination are experience of the past, then it's worth considering whether there are any cases of hallucination left over which involve experience of what doesn't exist. If not, then the argument from hallucination goes wrong in just the way I'm suggesting: by incorrectly assuming that hallucinations are experience of what doesn't exist.

Although a thorough discussion of the issues is outside the scope of this paper, it's worth sketching them. There are two sorts of cases which, even granting my proposal here, potentially involve experience of what doesn't exist: (1) leftover realistic cases of hallucinations not covered above, such as potential dream experiences of truly novel objects, and (2) highly speculative cases from philosophy in which there is no causal connection of any kind to physical objects which might be the objects of experience, such as in the case of a life-long fully envatted brain. The first cases don't provide any in-principle problem for the suggested response to the argument from hallucination, and might yet be explainable as experience of the past. It's an empirical question whether these cases of hallucination involve truly novel objects (i.e., of objects which aren't amalgamations of past-perceived object parts). As noted above, there's also the possibility of explaining potential cases of such hallucinations as highly distorted illusory experiences of past-perceived objects. For example, if I dream a strange man who doesn't *look* to be merely an amalgamation of parts from objects I have perceived in the past, it might simply be that I'm experiencing an object I did perceive in the past, but my experience misattributes to this object an unrecognizable form, shape, or set of properties.

The case which provides a real in-principle challenge to the approach on offer here is the second: the case in which the subject of the hallucinatory experience has never interacted, through their sensory systems, with any objects at all. Consider, for example, a life-long fully envatted brain. This

brain has never been in a body and has spent its time, from the start, in an artificial environment (its vat). It has, from the start, merely been fed input (through direct electromagnetic stimulation) based on computer simulations of a virtual (nonphysical) environment. In such a case it's simply not a possibility that the experiences of this brain are experiences of objects with which the brain has interacted in the past through its sensory systems, as this brain has never had any sensory interactions at all with physical objects. Not, anyway, if sensory interactions of a sort relevant to sensory experience are to be construed as fundamentally physical interactions: e.g., involving a body, sensory organs, and the transduction of energy between environment and organism through those organs.

If the main concern (for addressing the argument from hallucination) is to find candidate objects with which such an envatted brain *physically interacts*, then Raleigh provides a promising suggestion. He suggests (2014, 95) that the objects of these sorts of hallucinations are 'some element or structure within the hallucination producing apparatus'. While a life-long fully envatted brain has never physically interacted with traditional sensory stimuli through sensory systems, it does continually interact physically with the instruments, computer, and other equipment which feed it information about its virtual environment. So if the main concern is to avoid the argument from hallucination by finding some physical object, with which hallucinating subjects physically interact, which can be identified as the object of the hallucinatory experience, then Raleigh's approach offers one way to handle the case of life-long fully envatted brains. Of course, this approach doesn't make their hallucinations experience of the past, nor does it make the relevant interactions *sensory* (in my sense, as this involves the mediation through sensory organs). Still, my concern is not to show that *all* hallucination is experience of the past. The present aim is to show that the traditional argument from hallucination can be seen as going wrong precisely in assuming that the objects of hallucinations do not exist.

As Raleigh notes (2014, 103), his suggestion won't work for another case. Imagine a brain somehow forms, by random chance through quantum fluctuations or some such means, in a state that's causally disconnected from any physical objects, and also just happens (again by random chance) to have neural activity which perfectly matches the sensory activity in a normal subject. This brain (a "Boltzmann brain", or a "cosmic swampbrain", as Papineau calls it) would (so the intuition goes) have just the same sensory experience as the normal subject. But since this brain would be (and have al-

ways been) causally disconnected from the world, there would be no physical object with which it was sensorily interacting—not even the 'hallucination producing apparatus' (Raleigh 2014, 95) of the life-long fully envatted brain.

As Raleigh discusses (2014, 106), this case and similar ones (e.g., swampman) pose a challenge to representationalists about sensory experience who endorse informational or teleosemantic approaches to representation (e.g. Tye 1998; Lycan 2001; Dretske 2003). The standard move by these representationalists is to bite the bullet and say that the cosmic swampbrain lacks any sensory experience at all (and hence does not afford a case of hallucinating what doesn't exist). Raleigh notes (ibid) that Papineau has suggested that our intuitions about such a case aren't reliable anyway. Someone wanting to respond to the argument from hallucination by denying that hallucination ever involves experience of what doesn't exist should adopt this strategy, at least for the cosmic swampbrain, and perhaps (as an alternative to Raleigh's proposal) for the life-long fully envatted brain as well. An opponent who wants to press the traditional argument from hallucination through the cases of a cosmic swampbrain or life-long fully envatted brain goes wrong by putting weight on their (unreliable) intuition that such brains would have sensory experiences of objects. Note that we need not go as far as the representationalists and deny that these brains would have *any* experience. We only need to deny that they would have experiences *of objects*. Perhaps they would still be in some sort of phenomenal state, just not an intentionally directed state presenting an object. The upshot is that even if we can't avoid the traditional argument from hallucination by holding that all hallucinations are experience of the past, the insight that the argument goes wrong by assuming that the objects of hallucination don't exist can still be preserved.

Let me conclude by noting a final implication for the naïve realist. If it's true that there are no remaining cases of hallucinating *what doesn't exist*, then it's possible for naïve realists to account for both perception and hallucination as a relation to the object of experience. This is one of Raleigh's main points (2014), and the motivation for his proposal about envatted brains. While it's generally thought naïve realists must give some other, non-relational account of hallucination, i.e. that they must be disjunctivists of a certain sort (e.g., Smith 2002, 43–44; Martin 2004, 38–43; Crane 2006, 139; Logue 2014, 209; Genone 2016, 7; Moran 2019, 10), some naïve realists have tried to explain sensory experience in both perception and hallucination as a relation.[21] Johnston (2004) suggests that in hallucination we're aware of uninstantiated

property clusters. Knight (2013) suggests we're aware of nonexistent objects. But these attempts (unlike Raleigh's) still buy into the idea that hallucinatory experience cannot be a relation to physical objects. If hallucination is experience of the past, the objects to which we're related in hallucination can be physical objects—albeit past-perceived ones. This solution may depend on the naïve realist embracing an eternalist view of time on which things in the past exist (see Moran 2019), but I don't think that's the most serious hurdle to overcome in pursuing this approach to naïve realism.

Acknowledgements

I would like to thank Mohan Matthen, Benj Hellie, Evan Westra, Elliot Carter, Mason Westfall, and Simon Wimmer for their helpful comments on this material. Giulia Martina and Alex Morgan both provided line-by-line written comments on an earlier draft of this paper, and much thanks is owed to them. Finally, two anonymous referees from this journal both provided insightful, deep comments which lead to important improvements, clarifications, and reframing at the revision stage.

Contact

University of Toronto at Mississauga
The Network for Sensory Research
Department of Philosophy
Maanjiwe Nendamowinan, 6th floor
3359 Mississauga Road
Mississauga, Ontario L5L 1C6
michael.barkasi@gmail.com
ORCID: 0000-0003-4070-8196

Notes

[1]The conclusion only follows if we assume the relevant hallucinations from P1 seem to their subjects to be experiences of objects, but this is a point I'll assume in what follows. There are cases of what might be thought of hallucinations which don't present *objects*. For example, when viewing the Herman Grid Illusion one "sees" color patches that don't exist, and pressure to the eye produces bright "spots" or "stars". In these cases one (we

might be tempted to say) hallucinates color and shape instances, but these hallucinated items don't strike one as an object. The cases I'm concerned with are unlike these and involve hallucinatory experiences which seem to present their subject with an object.

[2]We need not assume, with Millikan, that the signals produced by these sensory systems are representations.

[3]Below I'll use the term 'perceptual experience', or sometimes to be more explicit, 'successful perceptual experience', for those sensory experiences in which a distal sensory stimulus (like an object) is successfully perceived. Perceptual experiences, on this use of the term, are just all those sensory experiences which are not hallucinations. Note that I'm not using the terms 'sensory' and 'perception' as psychologists often do, to distinguish between "raw" sensory stimulation and our "interpretations" of those stimulations.

[4]These ideas go back even to causal sense-datum theories (e.g. Russell 1912; Grice 1961). Even a causal sense-datum theorist could accept the claim that if an object-presenting experience is generated by a sensory interaction with an object O, then that experience is an *indirect* experience of O. The point is that if the experience is generated by a sensory interaction with O, then (on this view) O will be the cause of the experience in the right sort of way for it to be an indirect experience of O. This weaker indirect version of premise 2 isn't sufficient for my argument, but the point is just that the underlying idea that the object of experience is deeply related to causal structure (and hence sensory interaction) goes back a long way.

[5]How this tuning of synaptic weight works depends on things like neurotransmitter levels, cell type, structural position within a circuit, and type and position of the synapse relative to the neuron itself (Feldman 2012, 557–62). As proposed by Hebb (1949), one way some synapses get tuned, some of the time, is through correlated spiking between two connected neurons which strengthens the synapses between them (Feldman 2012, 556). High (>30Hz) and low spike rates (<10Hz), spike timing (e.g., one neuron firing within about 20ms of another), and coordinated spiking (e.g., multiple presynaptic neurons spiking synchronously) all also shape synaptic weights (Feldman 2012, 558). Depending on these factors spike propagation tunes synapses: strengthening some (long-term potentiation) and weakening others (long-term depression). Evidence suggests that this activity-dependent plasticity shapes neural circuits involved in sensory processing, including both low-level feature extraction and high-level object recognition in animal (including human) vision (Albright 2012, 230; Feldman 2012, 564).

[6]I have not reproduced a figure of Kanizsa's triangles, but I assume the reader is familiar with these images.

[7]Recall from footnote 1 that the argument only works if the relevant hallucinations seem to their subjects to be experiences of objects.

[8]As is well-known, sensory stimulation almost always, if not always, underdetermines its distal source (e.g., Fodor and Pylyshyn 1981, 173). But as noted by Albright (and also Tang et al. 2018), not every bit of receptor input requires completion by top-down processing. In this way Albright's view is definitely not a version of the predictive processing approach (e.g. Clark 2013), since Albright sees top-down signals not as an attempt to predict the incoming sensory signal, but instead as a way to fill-in defective sensory signals.

[9]It might have been objected that Kanizsa triangles aren't hallucinated *objects*, i.e. that

the relevant hallucinatory experience doesn't seem to present an object. I take it that this example of the vanishing ball is more clearly an hallucinated object: the hallucinatory experience seems to present an object (the ball) to the subject.

[10]I want to thank two separate anonymous referees for raising this objection and pushing me to clarify how my view intersects with approaches that treat hallucinations and dreams as a form of imagination.

[11]Also note that there's nothing which precludes imaginative states from affording awareness of the past. Indeed, although there's no room here to pursue this point, the considerations discussed here suggest a parallel argument for the claim that at least some imaginative experiences (e.g., imagining past events or counterfactual situations) involve awareness of the past-perceived objects memory traces of which form the basis of those experiences. Something like premises 4 and 5 hold for episodic memory. So we get the conclusion that episodic memory recall and those states based on it is a sensory interaction with the past. By applying an analog of premise 2 we would get that these imaginative states afford awareness of the past. Byrne (2010, 21) makes a similar point when he suggests that episodic memory preserves 'cognitive contact' with past-perceived things. By this he's following Campbell (2002) and is referring to *whatever* link perception provides which affords us the ability to talk and think demonstratively about objects, but the spirit is in line with what I'm arguing for here. Similarly, Debus (2008) has proposed a relational account of memory, akin to naïve realist accounts of perceptual experience. Again there's an obvious parallel between Debus's relational account of memory—at least, her claim that episodic recollection is a relation (2008, 406)—and this implication that episodic memory recall is a sensory interaction with the past.

[12]And, further, given that receptive fields are naturally larger in higher visual areas, it would be hard to tell if the difference in activity was due to deafferented hyperexcitability or the presence of the face stimulus.

[13]For an overview of the empirical evidence supporting this interpretation of dreams as involving the recombination of fragmented episodic memories, see (Nielsen and Stenstrom 2005) and (Windt 2015, 546–50, 614–17).

[14]One suggestion is that this generating activity will be tied up with the function of sleep and dreams, and so we should look there. A number of different functions have been proposed for sleep and dreaming (for an overview, see Windt 2015, 604–9). There is good evidence that sleep plays some role in memory consolidation (Windt 2015, 604), which might lend some support to the idea that some form of recall is happening. A problem here is that memory consolidation could be a function of *sleep* without being a function of *dreaming*; further, the ability to dream can be lost without cognitive impairment (Windt 2015, 607).

[15]For more on this motor blockade, see (Hobson et al. 2000, 826) and (Hobson 2009, 809).

[16]According to Windt (2018, 2590), these bodily experiences involve awareness of your actual sleeping body, albeit illusory awareness.

[17]Windt is careful to say that not all dreams are driven by bodily self-sampling in the way just described. Bodily self-sampling instead drives a wide and interesting class of dreams. Still, for our purposes the view provides an interesting proposal for understanding the sensory activities involved in dreams.

[18]Going back to the question of whether my view makes these hallucinations a form of imagination, what I'm proposing here for dreams perhaps fits well with the idea that dreams are a form of imagination (e.g. Ichikawa 2009).

[19]This actually raises another potential reply to the problem of novel objects. If I dream a strange man, it may be that I'm experiencing some past-perceived object, but my experience attributes to this object many properties it doesn't have. In that case, what I experience may look like a strange man that I've never seen before, but may really be a radically distorted past-perceived object.

[20]While most philosophers reject this sort of sense-data theory, contemporary psychologists and neuroscientists often take for granted that the objects of sensory experience (so-called percepts or images) are mental entities in the head; this is the picture one finds explicitly articulated in at least some popular textbooks (e.g., Kandel et al. 2000, 412) and implicitly baked into many empirical research articles. Although these scientists don't articulate the philosopher's argument from hallucination, it seems clear they are motivated at least in part by its spirit.

[21]There are a variety of approaches among naïve realists who do accept disjunctivism. Martin (2004) says hallucinations are mental states we can't introspectively discriminate from perceptual experiences. Fish (2009) says hallucinations lack any phenomenology at all. Kennedy (2013) says hallucination can be explained in terms of representation. Logue suggests that hallucinations lack perceptual phenomenology while still having representational content (2012a, 183). Hellie takes something of a sense-data approach to hallucination, suggesting hallucination involves awareness of a 'mental figment' (2014, 255), although elsewhere he suggests hallucinations likely have different explanations (2013).

References

Addis, D. R., A. T. Wong, and D. L. Schacter (2007). Remembering the past and imagining the future: Common and distinct neural substrates during event construction and elaboration. *Neuropsychologia 45*, 1363–1377.

Albright, T. D. (2012). On the perception of probable things: Neural substrates of associative memory, imagery, and perception. *Neuron 74*, 227–245.

Anscombe, G. (1965). The intentionality of sensation: A grammatical feature. In R. J. Butler (Ed.), *Analytic Philosophy*, pp. 158–80. Blackwell.

Austin, J. L. (1962). *Sense and Sensibilia*. Oxford: Oxford University Press.

Ayer, A. J. (1956). *The Problem of Knowledge: An Enquiry into the Main Philosophical Problems that enter into the Theory of Knowledge*. London: Pelican Books.

Blumberg, M. S. (2010). Beyond dreams: Do sleep-related movements contribute to brain development? *Frontiers in Neurology 1*(140), 1–10.

Blumberg, M. S. (2015). Developing sensorimotor systems in our sleep. *Current Directions in Psychological Science 24*(1), 32–37.

Blumberg, M. S., C. M. Coleman, A. I. Gerth, and B. McMurray (2013). Spatiotemporal structure of REM sleep twitching reveals developmental origins of motor synergies. *Current Biology 23*(21), 2100–2109.

Brentano, F. (2015/1874). *Psychology from an Empirical Standpoint*. Routledge Classics. Abingdon: Routledge.

Brogaard, B. and D. E. Gatzia (2017). Unconscious imagination and the mental imagery debate. *Frontiers in Psychology 8*(799), 1–14.

Buckner, R. L. and D. C. Carroll (2007). Self-projection and the brain. *TRENDS in Cognitive Sciences 11*(2), 49–57.

Burge, T. (2005). Disjunctivism and perceptual psychology. *Philosophical Topics 33*(1), 1–78.

Burge, T. (2009). Perceptual objectivity. *Philosophical Review 118*(3), 285–324.

Burge, T. (2010). *Origins of Objectivity*. Oxford: Oxford University Press.

Byrne, A. (2010). Recollection, perception, imagination. *Philosophical Studies 148*, 15–26.

Campbell, J. (2002). *Reference and Consciousness*. Oxford: Oxford University Press.

Chalmers, D. J. (2006). Perception and the fall from Eden. In T. Szabo Gendler and J. Hawthorne (Eds.), *Perceptual Experience*, pp. 49–125. Oxford: Oxford University Press.

Clark, A. (2012). Dreaming the whole cat: Generative models, predictive processing, and the enactivist conception of perceptual experience. *Mind 121*(483), 753–771.

Clark, A. (2013). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences 36*(3), 181–253.

Crane, T. (2006). Is there a perceptual relation? In T. Szabo Gendler and J. Hawthorne (Eds.), *Perceptual Experience*, pp. 126–146. Oxford: Oxford University Press.

Crick, F. and C. Koch (1998). Consciousness and neuroscience. *Cerebral Cortex 8*, 97–107.

Dayan, P. and L. Abbott (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. Cambridge: The MIT Press.

De Brigard, F. (2014). Is memory for remembering? recollection as a form of episodic hypothetical thinking. *Synthese 191*, 155–185.

Debus, D. (2008). Experiencing the past: A relational account of recollective memory. *Dialectica 62*(4), 405–432.

Dickie, I. (2010). We are acquainted with ordinary things. In R. Jeshion (Ed.), *New Essays on Singular Thought*, pp. 213–245. Oxford: Oxford University Press.

Dretske, F. (1988). *Explaining Behavior: Reasons in a World of Causes*. Cambridge: The MIT Press.

Dretske, F. (1995). *Naturalizing the Mind*. Cambridge: The MIT Press.

Dretske, F. (2003). Experience as representation. *Philosophical Issues 13*(1), 67–82.

Feldman, D. E. (2012). The spike-timing dependence of plasticity. *Neuron 75*, 556–571.

ffytche, D. H. (2013). The hallucinating brain: Neurobiological insights into the nature of hallucinations. In F. Macpherson and D. Platchias (Eds.), *Hallucination: Philosophy and Psychology*, pp. 45–64. Cambridge: The MIT Press.

ffytche, D. H., R. J. Howard, M. J. Brammer, A. David, P. Woodruff, and S. Williams (1998). The anatomy of conscious vision: An fMRI study of visual hallucinations. *Nature Neuroscience 1*(8), 738–742.

Fish, W. (2009). *Perception, Hallucination, and Illusion*. Oxford: Oxford University Press.

Fodor, J. A. and Z. W. Pylyshyn (1981). How direct is visual perception?: Some reflections on Gibson's 'ecological approach'. *Cognition 9*, 139–196.

Fosse, M. J., R. Fosse, J. A. Hobson, and R. J. Stickgold (2003). Dreaming and episodic memory: A functional dissociation? *Journal of Cognitive Neuroscience 15*(1), 1–9.

Genone, J. (2016). Recent work on naïve realism. *American Philosophical Quarterly 53*(1), 1–26.

Grice, H. P. (1961). The causal theory of perception. *Proceedings of the Aristotelian Society 121*, 121–152.

Hebb, D. O. (1949). *The Organization of Behavior: A Neuropsychological Theory*. New York: Wiley.

Hellie, B. (2013). The multidisjunctive conception of hallucination. In F. Macpherson and D. Platchias (Eds.), *Hallucination: Philosophy and Psychology*, pp. 149–174. Cambridge: The MIT Press.

Hellie, B. (2014). Love in the time of cholera. In B. Brogaard (Ed.), *Does Perception Have Content?*, pp. 241–261. Oxford: Oxford University Press.

Hinton, J. M. (1967). Visual experiences. *Mind 76*, 217–227.

Hobson, J. A. (2009). REM sleep and dreaming: Towards a theory of protoconsciousness. *Nature Reviews Neuroscience 10*(11), 803–813.

Hobson, J. A. and K. J. Friston (2012). Waking and dreaming consciousness: Neurobiological and functional considerations. *Progress in Neurobiology 98*, 82–98.

Hobson, J. A. and K. J. Friston (2014). Consciousness, dreams, and inference: The cartesian theatre revisited. *Journal of Consciousness Studies 21*(1-2), 6–32.

Hobson, J. A., C. C. H. Hong, and K. J. Friston (2014). Virtual reality and consciousness inference in dreaming. *Frontiers in Psychology 5*(1133), 1–18.

Hobson, J. A., E. F. Pace-Schott, and R. Strickgold (2000). Dreaming and the brain: Toward a cognitive neuroscience of conscious states. *Behavioral and Brain Sciences 23*(6), 793–842.

Hume, D. (1993). *An Enquiry Concerning Human Understanding* (2nd ed.). Indianapolis: Hackett Publishing Company.

Ichikawa, J. (2009). Dreaming and imagination. *Mind and Language 24*(1), 103–121.

Jackson, M. B. (2013). Recall of spatial patterns stored in a hippocampal slice by long-term potentiation. *Journal of Neurophysiology 110*, 2511–2519.

Johnston, M. (2004). The obscure object of hallucination. *Philosophical Studies 120*(1-3), 113–83.

Johnston, M. (2006). Better than mere knowledge? The function of sensory awareness. In T. Szabo Gendler and J. Hawthorne (Eds.), *Perceptual Experience*, pp. 260–290. Oxford: Oxford University Press.

Johnston, M. (2011). On a neglected epistemic virtue. *Philosophical Issues 21*(1), 165–218.

Kandel, E. R., J. H. Schwartz, and T. M. Jessell (Eds.) (2000). *Principles of Neural Science* (4th ed.). New York: McGraw-Hill Health Professions Division.

Kanizsa, G. (1976). Subjective contours. *Scientific American 234*(4), 48–53.

Kennedy, M. (2013). Explanation in good and bad experiential cases. In F. Macpherson and D. Platchias (Eds.), *Hallucination: Philosophy and Psychology*, pp. 221–54. Cambridge: The MIT Press.

Knight, G. (2013). Disjunctivism unmotivated. *Phenomenology and the Cognitive Sciences*, 1–18.

Liu, X., S. Ramirez, R. L. Redondo, and S. Tonegawa (2014). Identification and manipulation of memory engram cells. *Cold Spring Harbor Symposia*

on Quantitative Biology 79, 59–65.

Logue, H. (2012a). What should the naïve realist say about total hallucinations? *Philosophical Perspectives 26*, 173–199.

Logue, H. (2012b). Why naïve realism? *Proceedings of the Aristotelian Society 112*, 211–237.

Logue, H. (2014). Disjunctivism. In M. Matthen (Ed.), *The Oxford Handbook of Philosophy of Perception*, pp. 198–216. Oxford: Oxford University Press.

Lycan, W. G. (2001). The case for phenomenal externalism. *Philosophical Perspectives 15*, 17–35.

Macpherson, F. (2012). Cognitive penetration of colour experience: Rethinking the issue in light of an indirect mechanism. *Philosophy and Phenomenological Research 84*(1), 24–62.

Macpherson, F. (2013). The philosophy and psychology of hallucination: An introduction. In F. Macpherson and D. Platchias (Eds.), *Hallucination: Philosophy and Psychology*, pp. 1–38. Cambridge: The MIT Press.

Macpherson, F. (2014). The space of sensory modalities. In D. Stokes, M. Matthen, and S. Biggs (Eds.), *Perception and Its Modalities*, pp. 432–461. Oxford: Oxford University Press.

Martin, M. G. F. (2004). The limits of self-awareness. *Philosophical Studies 120*(1-3), 37–89.

Matthen, M. (2005). *Seeing, Doing, and Knowing: A Philosophical Theory of Sense Perception*. Oxford: Oxford University Press.

Matthen, M. (2014). How to be sure: Sensory exploration and empirical certainty. *Philosophy and Phenomenological Research 88*(1), 38–69.

Matthen, M. (2019). Objects, seeing, and object-seeing. *Synthese*, 1–24.

McDowell, J. (1986). Singular thought and the extent of inner space. In J. McDowell and P. Pettit (Eds.), *Subject, Thought, and Context*, pp. 137–168. Oxford: Oxford University Press.

Michaelian, K. (2016). *Mental Time Travel: Episodic Memory and Our Knowledge of the Personal Past*. Cambridge: The MIT Press.

Millikan, R. G. (1984). *Language, Thought and Other Biological Categories*. Cambridge: The MIT Press.

Millikan, R. G. (2004). *Varieties of Meaning: The 2002 Jean Nicod Lectures*. Cambridge: The MIT Press.

Moran, A. (2019). Naïve realism, seeing stars, and perceiving the past. *Pacific Philosophical Quarterly 100*(1), 202–232.

Nagel, T. (1974). What is it like to be a bat? *Philosophical Review 83*, 435–450.

Nielsen, T. A. and P. Stenstrom (2005). What are the memory sources of dreaming? *Nature 437*(7063), 1286–1289.

Noë, A. (2004). *Action in Perception*. Cambridge: The MIT Press.

Noë, A. (2007). Magic realism and the limits of intelligibility: What makes us conscious. *Philosophical Perspectives 21*, 457–474.

O'Callaghan, C. (2008). Seeing what you hear: Cross-modal illusions and perception. *Philosophical Issues 18*, 316–338.

O'Callaghan, C. (2016). Objects for multisensory perception. *Philosophical Studies 173*, 1269–1289.

Painter, D. R., M. F. Dwyer, M. R. Kamke, and J. B. Mattingley (2018). Stimulus-driven cortical hyperexcitability in individuals with Charles Bonnet hallucinations. *Current Biology 28*, 3475–3480.

Pautz, A. (2009). What are the contents of experiences? *The Philosophical Quarterly 59*(236), 483–507.

Pautz, A. (2010). Why explain visual experiences in terms of content. In B. Nanay (Ed.), *Perceiving the World*, pp. 254–309. Oxford: Oxford University Press.

Penfield, W. and P. Perot (1963). The brain's record of auditory and visual experience. *Brain 86*(4), 595–696.

Prinz, J. (2000). A neurofunctional theory of visual consciousness. *Consciousness and Cognition 9*(2), 243–259.

Prinz, J. (2006). Beyond appearances: The content of sensation and perception. In T. Szabo Gendler and J. Hawthorne (Eds.), *Perceptual Experience*, pp. 434–460. Oxford: Oxford University Press.

Prinz, J. (2011). Is attention necessary and sufficient for consciousness? In C. Mole, D. Smithies, and W. Wu (Eds.), *Attention: Philosophical & Psychological Essays*, pp. 174–203. Oxford: Oxford University Press.

Raleigh, T. (2014). A new approach to 'perfect' hallucinations. *Journal of Consciousness Studies 21*(11–12), 81–109.

Rosen, M. (2018). Enactive or inactive? Cranially envatted dream experience and the extended conscious mind. *Philosophical Explorations 21*(2), 295–318.

Russell, B. (1997/1912). *The Problems of Philosophy*. Oxford: Oxford University Press.

Schellenberg, S. (2010). The particularity and phenomenology of perceptual experience. *Philosophical Studies 149*, 19–48.

Schellenberg, S. (2011). Perceptual content defended. *Noûs 45*(4), 714–750.

Schellenberg, S. (2019). Perceptual consciousness as a mental activity.

Noûs 53(1), 114–133.

Schlack, A. and T. D. Albright (2007). Remembering visual motion: Neural correlates of associative plasticity and motion recall in cortical area MT. *Neuron 53*(6), 881–890.

Shams, L., Y. Kamitani, and S. Shimojo (2000). What you see is what you hear. *Nature 408*, 788.

Siegel, S. (2006). How does visual phenomenology constrain object-seeing? *Australasian Journal of Philosophy 84*(3), 429–441.

Siegel, S. (2010). *The Contents of Visual Experience*. Oxford: Oxford University Press.

Smith, A. D. (2002). *The Problem of Perception*. Cambridge: Harvard University Press.

Snowdon, P. F. (1980). Perception, vision, and causation. *Proceedings of the Aristotelian Society 81*, 175–92.

Soteriou, M. (2000). The particularity of visual perception. *European Journal of Philosophy 8*(2), 173–189.

Speaks, J. (2009). Transparency, intentionalism, and the nature of perceptual content. *Philosophy and Phenomenological Research 79*(3), 539–573.

Tang, H., M. Schrimpf, W. Lotter, C. Moerman, A. Paredes, J. O. Caro, W. Hardesty, D. Cox, and G. Kreiman (2018). Recurrent computations for visual pattern completion. *Proceedings of the National Academy of Sciences 115*(35), 8835–8840.

Tiippana, K. (2014). What is the McGurk effect? *Frontiers in Psychology 5*(725), 1–3.

Tulving, E. (1983). *Elements of Episodic Memory*. Oxford: Oxford University Press.

Tulving, E. (2002). Episodic memory: From mind to brain. *Annual Review of Psychology 53*, 1–25.

Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge: The MIT Press.

Tye, M. (1998). Inverted earth, swampman, and representationalism. *Philosophical Perspectives 12*(S12), 459–77.

Tye, M. (2009). Representationalist theories of consciousness. In B. P. McLaughlin, A. Beckermann, and S. Walter (Eds.), *The Oxford Handbook of Philosophy of Mind*, pp. 253–267. Oxford: Oxford University Press.

Valberg, J. J. (1992). *The Puzzle of Experience*. Oxford: Oxford University Press.

van de Ven, V. and A. T. Sack (2013). Transcanial magnetic stimulation of

visual cortex in memory: Cortical state, interference and reactivation of visual content in memory. *Behavioural Brain Research 236*, 67–77.

Windt, J. M. (2010). The immersive spatiotemporal hallucination model of dreaming. *Phenomenology and the Cognitive Sciences 9*, 295–316.

Windt, J. M. (2015). *Dreaming: A Conceptual Framework for Philosophy of Mind and Empirical Research*. Cambridge: The MIT Press.

Windt, J. M. (2018). Predictive brains, dreaming selves, sleeping bodies: How the analysis of dream movement can inform a theory of self- and world-simulation in dreams. *Synthese 195*, 2577–2625.