DOI: 10.1111/nous.12360

ARTICLE

Noûs

How to see invisible objects

Jessie Munton

Faculty of Philosophy, University of Cambridge, Cambridge, United Kingdom

Correspondence
Jessie Munton, Faculty of Philosophy, University of Cambridge, Sidgwick Avenue, Cambridge, CB3 9DA, United Kingdom.
Email: jm2200@cam.ac.uk

Abstract

It is an apparent truism about visual perception that we can see only what is visible to us. It is also frequently accepted that visual perception is dynamic: our visual experiences are extended through, and can evolve over time. I argue that taking the dynamism of visual experience seriously renders certain simplistic interpretations of the first claim, that a subject at a given time can see only what is visible to her at that time, false: we can be meaningfully said to see invisible objects. This counter-intuitive result in turn focuses our attention on the relationship between perception and memory. I show that it is difficult to draw a clear or simple distinction between the two. Memory and perception rely on, and blend with, one another. Together, these claims point us away from understanding visual perception as a simple reflection of the environment, and instead as closer to a process of dynamic modelling that draws together occurrent stimulation and stored information.

It is an apparent truism about visual perception that we can see only what is visible to us. It is also frequently accepted that visual perception is dynamic: our visual experiences are extended through time, allowing them to change and evolve through it. I argue that taking the dynamism of visual experience seriously renders certain simplistic interpretations of the first claim, that a subject at a given time can see only what is visible to her at that time, false. In particular, subjects can see objects during brief periods of occlusion. The primary relation a subject at a given moment

in time bears is to an extended period of visual experience, long enough to include moments at which certain objects may be briefly invisible. Once we prioritize this temporally extended experience over the constituent moments through which it extends, the subject can still be meaningfully said to see those objects despite their momentary invisibility in virtue of their visibility during the relevant extended period of experience. That is the first argumentative goal of this paper.

This is in some respects a highly counter-intuitive result. One way we might defend against it is to claim that the kind of extended, dynamic representations which support that claim are not genuinely perceptual. Instead, they are representations *in memory*. But when we try to clarify the requisite distinction between memory and perception, we find that it is hard to draw a clear line between the two. Instead, representations from memory are intimately woven into perceptual representations, and similar or identical neural processes support the two. Memory and perception are not opposed categories: perception, even of the present, relies on memory in a range of ways. This is the second argumentative goal of this paper. Between them, these two claims point us away from understanding visual perception as a simple reflection of the world around us, and instead as closer to a process of dynamically modelling the environment, by integrating occurrent stimulation and remembered information.

The paper proceeds as follows. In section one I set out in more detail the claim that visual experience is dynamic. In section two, I argue that allowing that the primary unit of visual object perception is temporally extended pushes us towards allowing that we can be meaningfully said to see invisible objects. In section three I consider the objection that such representations are not genuinely visual, and argue that it relies on an unsustainably rigid distinction between perception and memory.

1 | DYNAMIC PRIORITY

Philosophers commonly talk of having "a" visual experience - of a tomato, or a table, or a car crash - as though visual experience came in *units*. This invites the question: what is the temporal duration of such a unit? Does visual experience present the subject with a series of static snapshots with no temporal extension of their own, or does the most basic unit of visual experience develop and extend over time?

What would it mean to claim that visual experience is temporally extended?[1] It is clear enough that I am having a visual experience now, and was having one five minutes ago, and that there is some kind of continuity between those. Equally, one can, if one so desires, identify and talk of momentary snapshots within that period of experience. For claims about the temporal extension of visual experience to be non-trivial, they must be claims about metaphysical *priority*: are momentary snapshots the *basic* unit of visual experience, from which extended periods of experience are constructed, or is the most basic visual relation to an extended period of experience that presents something of temporal breadth, within which we can then pick out particular moments?

I will call the view that visual experience is fundamentally temporally extended, Dynamic Priority.

Dynamic Priority: the minimal "unit" of visual experience to which a subject is related is temporally extended

If Dynamic Priority holds, an extended period of experience is not determined by its successive instantaneous temporal parts.[2] From a representationalist perspective, one way of understanding Dynamic Priority is as the claim that the content and character of a subject's experience at

a moment in time is indeterminate: they depend on facts about a longer period of time within which the experience is embedded.

The important contrast here is with a view on which the subject is primarily related to instantaneous snapshots of experience, which could together comprise an experience that presented something of temporal breadth, in such a way that that extended experience would be secondary to, or metaphysically dependent on, those instantaneous snapshots.[3] Dynamic Priority is, as its name would suggest, a claim about the *priority* of a certain temporal scale of visual experience. It does not deny that we *can* individuate instantaneous snapshots of experience for our own theoretical purposes, only that such an isolation is secondary to, or parasitic on, temporally extended experience. This is the same kind of metaphysical priority that Ian Phillips describes in his (2011):

> What does it mean to say that a duration of experience is metaphysically prior to its sub-parts? It is not to deny that there are facts about instants during our stream of consciousness. It is, however, to insist that such facts are derivative. The most basic facts about our experiential lives are, in the first instance, facts about extended periods of the stream of consciousness. What is true at an instant is true only in virtue of that instant being an instant during a certain period of experience. (Phillips, 2011, 398)

A dynamic priority claim is a natural fit with the kind of extensionalism about temporal experience which Phillips endorses (Phillips, 2010), a view on which, as Barry Dainton puts it, "our episodes of experiencing are themselves temporally extended", in a way that allows our experience of temporal properties to rely on the temporal extension of the experience itself (Dainton, 2018).[4] By contrast, the retentionalist claims that our experience of temporal properties such as change can be generated at a moment, that it does not require that our experience itself be extended through time. On this view, experiences which themselves have no or minimal temporal duration in objective time can nonetheless provide a window onto an extended period of subjective time (G. Grush, 2005; Lee, 2014). Depending on the details of their position, some form of dynamic priority may also be endorsed by the retentionalist, either because they allow that the relevant atoms are themselves temporally extended in some significant way, or because they agree that at least certain experiences or experiential contents emerge from the pretension or retension of a series of such atoms, in a way that cannot be reduced to their individual contents. Dynamic priority is inconsistent only with the kind of "atomism" described by Philippe Chuard (2011), on which perceptual experience is a series of static snapshots with no presentation of temporal features.[5]

That literature on temporal experience is chiefly concerned with our experience of temporal properties themselves, such as change, or movement. By contrast, I am interested in a dynamic priority claim about visual perception of objects in particular – cases when a subject sees something *as* a spatio-temporally continuous particular. Why focus on object perception? Increasingly, work in vision science takes perceptual objects, rather than space per se, to be the fundamental currency of visual perception (Scholl, 2001; Vickery & Chun, 2010). This meshes with a philosophical emphasis on the attributive nature of visual perception (Burge, 2010) and its particularity (Schellenberg, 2016): both presuppose the centrality of objects to visual perception. If *object* perception is temporally extended, in the sense that dynamic units have priority over momentary time slices, then much of visual perception is inherently dynamic.

Applying Dynamic Priority to object perception, we arrive at the following claim:

Dynamic Priority (objects): the minimal "unit" of visual experience to which a subject is related when seeing an object is temporally extended.

What does this really mean? Imagine a subject, Bernadette, who is looking at a banana. Bernardette's experience of the banana extends from t_1 through to t_5. We can zone in on her experience at t_4. One way of understanding *Dynamic Priority* is as the claim that it wouldn't be *possible* for Bernadette to have whatever experience she is having of the banana at t_4 without the experience that occurs between t_1 and t_5 of which t_4 is a part. The moment is parasitic on the whole. If visual experience is genuinely dynamic, we should think of an individual at a moment in time as primarily related not to a momentary slice of visual experience that corresponds to the instant under investigation, but to an extended period of subjective experience which includes their experience at that moment.[6] A representationalist could say that there is no fact of the matter what the content or character of Bernadette's experience of the banana is at t_4, *except* facts that are derived from her temporally extended experience. Her experience at t_4 does not supervene just on facts about her experience at that time, but on facts about some longer period of time.

Here is another way of understanding the claim of interest. Suppose we could successfully specify what Bernadette is seeing at t_4 without referencing or knowing anything about her experience at t_1-t_3 or at t_5. In such a case, is Bernadette plausibly seeing an object at t_4? Dynamic object perception claims that the answer to this question is 'no': we can only attribute genuine object perception to a subject whose visual experience has some minimal temporal extension, and if we can fully specify Bernadette's visual experience at a moment, then whatever its content it cannot include the visual perception of an object.

Dynamic object perception leaves open what Bernadette *does* see in this case: it is plausible that certain properties could be momentarily perceptible even if objects were not (colour, contrast, or illumination, for instance that need not be bound to a visual object). But dynamic object perception could also be coupled with, or taken to naturally support, additional dynamic priority claims, for instance that certain properties which may be integral to object representation (such as size or shape) also require temporally extended visual experience.

Auditory experience is often thought to be temporally extended in this way (Nudds, 2010; O'Callaghan, 2009). Consider listening to a cadence at the end of a phrase of music. A phrase of music ending in a perfect cadence and one ending in an interrupted cadence may conclude on the very same chord. It can seem as though one's experience of hearing that final chord differs in the two cases not just in the sense of completion or homecoming that the perfect cadence has, and which the interrupted cadence promises and then turns away from, but in the more basic auditory experience of hearing the notes themselves.[7] One hears a middle C *as the concluding note* of a perfect cadence, for instance, rather than the concluding note of an interrupted cadence. Dynamic priority can naturally accommodate this holistic element to our auditory experience: some longer period of auditory experience has priority over the moment, whose development and internal structure and relations partly determine the experience picked out at any instant within that period of time. According to this interpretation of the cadence case, facts about the subject's auditory experience at t_5 depend on facts about their experience through moments t_1-t_5. If Dynamic Priority about object perception is right, then visual experience of objects is thoroughgoingly like our auditory experience of cadences in this respect: the subject's visual experience is more than the sum of its parts.

2 | FROM DYNAMIC PRIORITY TO SEEING INVISIBLE OBJECTS

A wide range of motivations could move one to endorse Dynamic Priority about object perception. For some, it will seem to be a necessary truth about visual object perception, one which falls out of

conceptual considerations of what it is to see an object.[8] For others, it may seem to have the status of a truth they can read off from their own visual phenomenology. For still others, phenomena such as visual constancy or our perception of motion might seem to require that something in the vicinity is true.[9] And finally, the strongest support for the claim may arise from some combination of these various considerations.

My concern here is not to weigh the evidence in its favor, but to consider what follows from it. According to Dynamic Priority (objects), the primary unit of visual experience involved in object perception is temporally extended. The content and character of the experience at a moment in time taken on its own is indeterminate: it is only when it forms part of a longer period of experience that it comes to be a moment of object perception. What ramifications does this have for what we can be said to see? I will argue that a dynamic priority claim for object perception disrupts our ordinary understanding of what we see because it loosens the reins between the input the visual system receives *at a moment* and the subject's conscious visual experience at (or very slightly after) that time. This loosening opens the way to allowing that we can see objects even during brief periods of invisibility.

Consequently, endorsing Dynamic Priority for object perception involves denying

Occurrent Visibility: A subject S sees an object o at time t only if o is visible to S at t_*.

where t is offset from t_* by just so long as it takes for the visual system to perform the relevant visual processing, that is, probably within a few hundred milliseconds at the most, ((VanRullen & Thorpe, 2001)).

Put another way, endorsing Dynamic Priority in the context of object perception entails the following claim:

Seeing the Invisible: a subject may see objects at time t which are not visible to her at t_*.

What is built into the notion of visibility and invisibility in play in this claim? Since we're interrogating the relationship between seeing and invisibility, we must work with a definition of the latter that is independent of the former. An object is invisible, for my purposes, when it is incapable of reflecting light in a manner that could be detected by a light sensitive surface such as the retina, or when, though the relevant light patterns are detectable by retinal photoreceptors, that information cannot be transmitted to the cortical visual system. Objects are most straightforwardly invisible in this sense when they are occluded, or alternatively when some block on the visual system prevents their uptake: you are invisible to me if I have my eyes shut.[10] Our ordinary notion of invisibility is indexed fairly precisely to a time: invisible ink is invisible on the page though it was visible when we wrote with it. Harry Potter is invisible when under his invisibility cloak, but visible without it. This time-indexing gives rise to the tension I am interested in, with a dynamic view of object perception.

Our focus for the moment then is on the question of whether a subject can be said to *see* an object, though it is temporarily incapable of reflecting light in a way that could be detected by the subject. Our ordinary conception of sight is tightly tied to that ability. I shall argue that Dynamic Priority puts pressure on that, encouraging us to reconceive of visual perception as a process of modelling both visible and invisible elements of a spatially and temporally extended scene.

To proceed with this argument, we first need a criterion of seeing. Our concern at the moment is not to arrive at a full specification of such a criterion, but to explore whether the invisibility of an object at a brief moment in time automatically disqualifies the subject from being said to see

the object. Unfortunately, there is little agreement about what it takes to see an object, but there is at least some consensus that seeing involves a phenomenal experience of a particular kind, which must bear an appropriate causal relation to the object in question, so that, for instance, the subject's conscious experience represents the appearance of the object with some reasonable degree of sensitivity.

Seeing an object: A subject S sees an object O when they (i) have a conscious visual experience of O that is (ii) reasonably sensitive to O's appearance in virtue of (iii) an appropriate causal relation to O.

This is intended to be schematic. The key question we are pursuing is whether momentary invisibility is inconsistent with a reasonable precisification of these three key elements once a dynamic priority claim is assumed. I will argue that it is not, and that we should consequently allow that the subject sees the object throughout the brief period of invisibility.

2.1 | Saccades

To see how Dynamic Priority supports these claims, consider what happens during moments when the flow of retinal information is interrupted, for instance when you saccade, or blink. A saccade occurs when you refocus your attention within a visual field by skipping both eyes across it simultaneously. Doing this repeatedly lets you build up a more accurate representation of the environment by moving the fovea, the most light-sensitive part of the retina, to acquire high quality information about a range of points within the visual field (Hollingworth, Richard, & Luck, 2008). Though these jumps typically happen at least a couple of times a second, we do not generally notice them occurring. To achieve this apparently seamless continuity, your sensory input partially cuts out whilst the movement occurs (Vallines & Greenlee, 2006). Otherwise, you would have a blurred interlude in the experience, as though viewing the world from a fast-moving vehicle.

Deliberate manipulation in experimental conditions reveals a period of desensitization that begins somewhere between 75 and 50 ms prior to the saccade, peaking at its start, and persisting through it. (Krekelberg, 2010; Ross, Morrone, Goldberg, & Burr, 2001). During this time, "stimulus-driven responses in early areas of the visual pathway are selectively inhibited" (Greenlee & Kimmig, 2019, 175). The diminution of sensitivity is sufficiently dramatic during this period of "saccadic suppression" that experimental subjects may fail to detect spatial shifts within a certain scale in an object's location if the object moves whilst the eyes are in motion (Li & Matin, 1990).[11] Even outside of experimental conditions, the impossibility of catching oneself saccading in a mirror is a tell of this temporary reduction in stimulus sensitivity (Krekelberg, 2010).

During the period of saccadic suppression, the object cannot be detected by the visual system and is in this sense invisible. Nonetheless the subject *sees* the object during this time period. Measuring the subject's experience during this time against our criterion of seeing we find that the subject has a conscious visual experience of the object that is reasonably sensitive to the object's appearance: it accurately represents its key featural properties in spite of the limited diminution in sensitivity to location.[12] The subject has that visual experience in virtue of a causal relationship with the relevant object. There is a brief temporal displacement between the cause and the conscious experience of it, but outside of carefully constructed experimental conditions this is not sufficient to disrupt the sensitivity to the appearance of the stimulus described above. At this level

at least, momentary invisibility seems perfectly compatible with seeing. The case of saccades is significant because it reveals that the functioning of the visual system is designed around the fact that objects are not continuously visible to it: it is constantly compensating for diminutions in the flow of information from the retina.

It might be objected that the object is not straightforwardly invisible when the subject saccades – others present could still see it after all. But we can easily construct a parallel hypothetical case in which saccadic suppression is instead paralleled by intermittent invisibility, for instance, imagine an object that passes behind a series of occluders at just the right rate – say for roughly 50 ms at a time, depending on lighting conditions – such that the visual system is unable to detect its periodic absences, producing instead a perceptual experience as of a continuously present object.[13] (This seems mechanically feasible: we know, for instance, from various apparent motion effects that we can have an experience as of an object in continuous motion, when briefly shown stimuli separated by intervals of between 25 and 400 ms (Colman, 2009)). In this intermittent occlusion case the subject would have a conscious experience of the object which is sensitive to its appearance, in virtue of a causal relationship with the object in question. The sensitivity is not *so* fine-grained in these cases that the subject detects the moments of invisibility, and this is where the Dynamic Priority thesis plays a role. The Dynamic Priority thesis for objects claims that object perception is temporally extended, such that the content and character of the subject's experience at those brief moments of invisibility is determined by the content and character over a longer period of time. Over those several seconds the subject's experience *is* sensitive to the object's appearance, and it is entirely natural to say that the subject sees the object. In light of that, when we assign a content to the moment of saccadic suppression or hypothetical disappearance, the subject can be said to see the object in spite of that momentary invisibility or suppression of neural activity.

The take home from both these kinds of cases is that in line with Dynamic Priority (objects) there is temporal flexibility in the causal relationship demanded by our ordinary notion of seeing, such that the subject sees the object despite its momentary invisiblity. But how far does that extend? In the cases above, the subject is not *conscious* of the object's invisibility. Could a subject see an object that is momentarily invisible in cases when that invisibility is a feature of their conscious experience? I turn next to considering certain cases of occlusion which, I shall claim, meet this specification.

2.2 | Occlusion

Momentary occlusion occurs when one object passes briefly behind another. When observing occlusion of this kind, you seem to *see* the object disappear and emerge from behind the occluder as part of a single coherent experience of its uninterrupted trajectory. You do not simply experience an object approach the occluder and incrementally disappear, and then a distinct object incrementally appear on the other side of it, though the latter event would involve an identical pattern of visual stimulation. This phenomenon of seeing the object as having a continuous trajectory whilst passing behind the occluder and out the other side is called the tunnel effect (Burke, 1952). Figure 1, provides an illustration of this kind of case, where the episode takes place over a few seconds.

Compare this with a case in which the object implodes just as it approaches the occluder and then explodes on the other side, that is, it shrinks to nothing and then suddenly grows from nothing on the other side. In this case, you have the experience of seeing the object approach the occluder, disappear and then reappear. Cases like this are designed to be as statistically similar

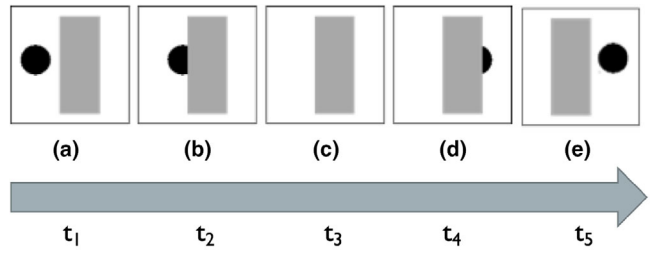

FIGURE 1 An object passing behind an occluder and emerging on the other side of it [Color figure can be viewed at wileyonlinelibrary.com]

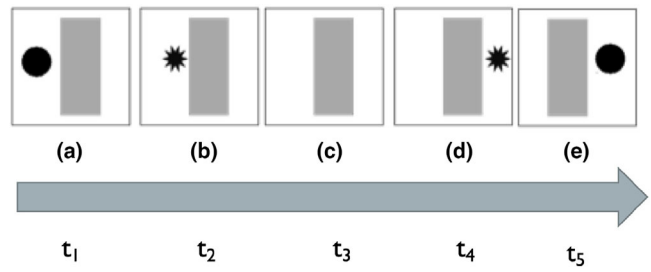

FIGURE 2 An object approaching an occluder, imploding, and then another object exploding on the other side [Color figure can be viewed at wileyonlinelibrary.com]

to the occlusion case as possible, whilst not giving rise to the tunnel effect. Figure 2 provides a schematic illustration of such a case.[14]

What do you see at t_5 in the two cases above? The same stimulus is available to the eye in both cases at that moment: a round, black, object on the right-hand side of the rectangular occluder. But what observers are prone to report seeing is different in the two cases. In the occlusion case, one sees the object as having emerged from behind an occluder, whereas in the implosion case one sees the object as having just appeared beside the occluder. As when listening to a cadence or an utterance, what the subject experiences at t_5, when the object is on the other side of the occluder, is palpably *different* in virtue of that moment in time being a part of an extended period of visual experience, that includes the preceding moments during which the object moved behind the occluder. The content and phenomenology of visual experience at a moment seems to depend on the longer period of time of which that moment forms a part.

We can shore up this self-report of the phenomenology of seeing an object pass behind and emerge from an occluder with some more objective investigation of our visual representations during occlusion. Scholl and Pylyshyn, (1999) have investigated these kinds of contrast cases using an object tracking paradigm. Subjects had the task of tracking a number of target objects, which either passed behind an occluder, or which appeared to implode on one side of it, and explode on the other. They found that subjects are better at tracking objects that disappear and reappear than those that implode and explode. Similarly, it is harder to track featural changes (such as color) of objects that implode or explode than of objects that are occluded. Change detection is better when it occurs within the same persisting object representation (Flombaum & Scholl, 2006).

Considering the representational vehicles posited to underlie visual object perception helps us understand these effects. Object perception relies on so-called "object files", a kind of mental folder that is opened when something is identified as an object, as a store for information about it. (Green & Quilty-Dunn, 2017). Objects files are temporally extended: they persist through feature-change, allowing you to see an object that changes color as it moves, for instance, as just that, rather than as a series of distinct objects, of different colors, at different locations.[15] Object files let us identify objects *as the same object* despite changes in spatiotemporal and featural properties, such as shape, colour and category (Kahneman, Treisman, & Gibbs, 1992; Scholl & Flombaum,

2010). The fact that that information is assimilated to one and the same persisting object file is what gives rise to a representation of change, or movement, consistent with the survival of the same object.

Appeal to object files helps us understand the differences between our experience of briefly occluded moving objects, as compared with an imploding/exploding pair of objects. One explanation for our superior abilities at identifying feature change in the former is that object representations are maintained during occlusion, but not during statistically equivalent events inconsistent with the on-going survival of the object in question. The apparent implosion of an object closes the relevant object file. Another file is opened when the object explodes into existence. The change in object files makes it much harder to track feature change in the relevant items. (Flombaum & Scholl, 2006; Scholl & Pylyshyn, 1999). This is consistent too with evidence of neural activity within the visual cortex which continues despite the occlusion of the object in question, suggesting the active maintenance of the relevant object representation (Hulme & Zeki, 2007).[16]

To recap: when an object passes behind an occluder and emerges again, both the phenomenology of seeing that object, and the kinds of tasks we're able to perform in relation to it, are different from a statistically closely equivalent case in which an object approaches an occluder, disappears, and then an object reappears on the other side of the occluder, in a manner inconsistent with its survival. The object file survives occlusion but not implosion, and that difference has an impact on the content and character of the subject's experience at the point of invisibility. During occlusion the object is experienced *as occluded* whilst in the case of implosion it is not, and that difference at t_3 makes a difference to the subject's visual experience of the object at t_5, when it reappears, in a manner consistent with Dynamic Priority. The persistence of the underlying object file provides a mechanism for the integration of information over an extended time period, of the kind that Dynamic Priority requires.

Note that object files are neither necessary nor sufficient for seeing an object: we may see objects without object files, as when they feature in the background of a scene but we fail to attend to them, or when we see a larger number of objects than we could maintain separate object files for,[17] and there is some evidence that we can maintain an object file that fails to correspond to an experiential object (Mitroff, Scholl, and Wynn (2005)). Object files are significant for our purposes because they reveal how the apparatus of conscious object perception is built around a temporally extended mechanism that allows for at least the subpersonal maintenance of information specifically about an object's appearance even during brief periods of invisibility in a manner that also seems to impact on a subject's phenomenal experience.

2.3 | What does the subject see?

What should we say about the subject's visual experience at t_3 in the occlusion case, the point at which the object is behind the occluder? What can the subject be said to *see* at that moment in time, if we treat the extended period of experience as the primary unit of visual perception?

The experimental evidence suggests that the subject is maintaining the relevant object file throughout the period of time. But our criterion of seeing requires not just the presence of a subpersonal representation of this kind, but a *conscious* visual representation of the object. The subject's experience over the whole period of time is *of the object passing behind the occluder* and so at t_3 the subject's experience is *of the object behind the occluder* that is, the content and character of experience is not just sensitive to the occluder, but *to the object behind it as well*: that is what gives rise to the tunnel effect. The object continues to play an active role in determining the subject's phenomenology even whilst invisible. If we want to capture the difference in the conscious

phenomenal experience of seeing an object pass behind the occluder and seeing it implode and explode next to the occluder, we need to appeal to the fact that, for this brief period of time, the subject represents the object as it passes behind the occluder, *despite its temporary invisibility*.

If that is right, then we sometimes see objects even when they are not visible to us: at t_3 a subject S may see an object O though at t_3 O is not visible to S, in virtue of the fact that t_3 is part of a longer time period (t_1-t_5) throughout which S maintains a visual representation of O, which is causally related to O in a straightforward way. In this way, accepting Dynamic Priority should lead us to accept

Seeing the Invisible: a subject may see objects at time *t* which are not visible to her at $t_{(*)}$.

It is worth drawing out that Seeing the Invisible is inconsistent with

Occurrent Visibility: A subject *S* sees an object *o* at time *t* only if *o* is visible to *S* at $t_{(*)}$.

Taking seriously the dynamic nature of visual object perception requires us to relax our focus on *momentary* invisibility, and take a longer view, so to speak, of the relevant causal relationship between the subject and their environment.

2.4 | Tightening up

I have made a case that the subject's relation to an occluded object can satisfy the first element in the schematic criterion of seeing outlined above: they have a conscious visual experience of the object. But what about the second and third – the requirement of reasonable sensitivity to the object's appearance, in virtue of an appropriate causal relationship to the object. I claim that this case also satisfies a reasonable interpretation of each of these.

This claim faces some obvious problems in the case of an object that is invisible: *how* can a subject bear an appropriate causal relationship to an object when momentary invisibility interrupts that causal relationship? And how can a causal chain that is interrupted in that way facilitate a reasonable degree of sensitivity to the object's appearance, given that the object is concealed behind the occluder at that time?

Consider first of all whether the conscious experience bears an appropriate causal relationship to the object. During occlusion, there are two causal relations in place. The presence of the object before and, crucially, during the onset of occlusion causes the on-going representation of it during occlusion. This causal relationship is temporally displaced to some degree: the causal relationship between the object at t_{1*}-t_{2*} is responsible for the subject's visual experience of the object at t_1-t_2, but it plays a significant role in sustaining the experience during occlusion, that is at t_3 and beyond, until the object emerges from behind the occluder again. In particular, the gradual way in which the object disappears behind the occluder allows the subject to continue to maintain the relevant object file in a way that implosion does not. And so in addition, this delayed causal relationship is coupled with an *on-going* causal relationship between the subject and the occluder. That sustains the representation of the occluder, but also the representation of the object *as occluded*. In section 3 we shall explore in more depth whether the temporal displacement in the first of those causal relationships is consistent with visual perception. For the moment, note that this is only a slight exaggeration of the displacement in the saccades case. Cases of saccading and blinking, and the extension of object files themselves all suggest that the causal

relationship between stimulus and conscious experience allows for some degree of temporal flexibility.

Secondly, we need to consider whether such a displaced temporal relationship can support a reasonable level of sensitivity to the object's appearance during the moment of occlusion. Even granting that the subject has a phenomenal awareness of the object, is that awareness reasonably sensitive to the object's appearance? How can a causal relationship support sensitivity to appearance at a later moment, particularly when it has already sustained a visual experience that is more temporally proximate to the cause? This difficulty manifests in the way in which the experience of the object during occlusion lacks the vivid phenomenology and the featural detail characteristic of visual representation of a visible object: if the object changes shape or color during its occlusion, our visual representation of it at that time will not be sensitive to that.

The subject is sensitive to the appearance of the object in two respects. Firstly, the subject represents the object *as occluded* and that is reflected in the visual appearance of the object which is obstructed by the presence of the occluder. It might be objected that this demonstrates that we are sensitive to the appearance of the occluder, not the object itself, and so it is worth noticing that in addition to this, the subject *does* retain information about the appearance of the object itself. Object files allow for residual sensitivity to phenomenal features despite the object's temporary invisibility. Consider the finding that feature change was more readily detected in cases of occlusion than implosion (Flombaum & Scholl, 2006). That suggests that the subject *does* continue to represent the object's features during occlusion at least subpersonally, not just whether or not the object continues to exist.

The conscious detail or vividness of the presentation of those features is muted, of course. There is in this respect a significant phenomenological difference between seeing occluded and non-occluded objects. However, there are other cases too in which our perception of an object may occur at so low a grade of resolution that we are incapable of perceiving its features in detail, for instance when we see things at a distance or at the periphery of our visual field. An object at a distance may present with sufficient ambiguity that we refrain from ascribing a color or size to it. Nonetheless, it remains the case that we see the object, despite the absence of accurate featural perception. We should not set the standard for reasonable sensitivity to the object's appearance so high that we exclude the possibility of seeing objects at a distance.[18] The case of object perception during occlusion is an extreme case of this: the lack of featural detail during the period of occlusion is consistent with the visual representation of the object in question, particularly coupled with the way in which the subject veridically represents the object *as occluded*.

It is important to remember that we are presupposing Dynamic Priority (objects). This is significant because it shifts our frame: the moment is secondary to the longer period of experience. Consider one straightforward 'fix' that would disqualify us from being said to see invisible objects: simply building into the definition of what it is to see something that the object in question must be visible to the subject at the moment when sight of it is ascribed to them. That is, we exclude in principle possible vision of invisible objects.[19]

But what motivation do we have to adopt such an exclusion? The insistence that the subject cannot see the object in virtue of its invisibility depends on denying that priority of the extended visual experience. It assumes the priority of the snapshot, and allows that visibility during that moment is a criterion on what the subject can be said to see. But why should *momentary* (in)visibility be the benchmark for whether a subject sees an object, once we recognize that the moment is metaphysically subordinate to the extended experience? If Dynamic Priority is right, the subject never sees an object at a single moment in time: the nature and function of object perception points us towards understanding such representations as inherently temporally extended. Our understanding

of what a subject sees at a given time should then be similarly sensitive to the priority of the extended experience. The subject sees the object in virtue of its visibility during the majority of the experience, and the on-going visual representation of it during the period of invisibility.

Let's compare this case of momentary total occlusion with a case of partial occlusion across a longer period of time. Suppose that I see my bicycle standing on the other side of a fence. If you ask me if I can see my bicycle, it is clearly appropriate for me to answer yes, and not, for instance, to say "well I can see most of the front wheel and all of the back, and a few chunks of the frame itself" even though the parts of the bicycle that are occluded are not themselves causing my representation of the bicycle, and though the visual experience's sensitivity to their appearance is somewhat reduced (they could, for instance, change colour without my detecting it). We reach this verdict because we tend to treat entire objects as the primary *spatial* unit of perception.[20] Similarly, it naturally follows from Dynamic Priority that we can see an object despite its total occlusion at a moment in time, since the primary unit of experience is not thin temporal slices, but a longer period of experience which is capable of supporting a conscious representation of the object behind the occluder which is reasonably sensitive to the object's appearance, despite its momentary invisibility.

There is a risk that this argument gets lost in the semantics of what it takes to see something, when the real import is for how we think of visual perception and what its fundamental task is. We allow that you see your bicycle though the fence covers some portion of it, because we think of visual perception as a relation to an object that can survive these partial elements of occlusion. Appreciating the dynamic nature of object perception should push us towards a temporally extended understanding of object perception which can include temporal as well as spatial pockets of invisibility. In the next section I turn to considering the ramifications of this for our understanding of the boundary between perception and memory.

3 | PERCEPTION OR MEMORY?

The argument above relies on the claim that the visual experience the subject has during occlusion is caused by the occluded object. This requires the causal relationship between object and experience to accommodate a short time lag, in addition to that introduced by regular perceptual processing: the subject's experience of the object at t_3, when it is behind the occluder, is caused by the object at t_{1*}-t_{2*},. Could this ground an objection to the claim we see invisible objects? Causes always (perhaps necessarily) precede their effects, frequently by a prolonged period of time. But, it might be claimed, the close temporal proximity of cause and experience is what distinguishes genuine perception from memory. By restricting the relevant causal relation between object and experience to one that is as temporally proximate as visual processing allows, we can categorize experience of momentarily occluded objects not as a genuinely visual or even perceptual phenomenon, but rather as a process of cognitive representation, grounded heavily in memory. We cannot *see* invisible objects but we can *remember* them. Perception is in part distinguished from memory by its function of accurately reflecting the world currently before us, that is, the *visible* world in particular.[21]

When discussing perception of motion, Reid advances something like this position:

> [T]hough in common language we speak with perfect propriety and truth when we say that we see a body move, and that motion is an object of sense, yet when as Philosophers we distinguish accurately the province of sense from that of memory,

> we can no more see what is past, though but a moment ago, than we can remember what is present; so that philosophically speaking, it is only by the aid of memory that we discern motion, or any succession whatsoever. (1785/2015, -271)

Reid's position can be readily extended to object perception during occlusion: accurately distinguishing the provinces of sense from memory will reveal to us that we cannot *see* objects during occlusion. Whatever representations we maintain of them are in memory, not perception.

More recently, Elizabeth Spelke has advanced a similar line of argument, that object perception *per se* is not genuinely visual. One strand of Spelke's argument is particularly relevant here: her appeal to the fact that objects continue to be represented at times when they are not visible, for instance when they are fully occluded. In light of this representation during periods of invisibility, Spelke argues that object representation is not a genuinely visual phenomenon. Rather, "[t]he apprehension of objects is a cognitive act, brought about by a mechanism that begins to operate at the point where perception ends" (Spelke, 1988, 199).

Applying this approach to the case of occlusion, we could allow that you *represent* or *visually refer* to the object even when it is occluded, but insist that what you *see* exclude representations whose cause is in the past to some significant degree (more than the brief delay mandated by the need for perceptual processing).[22]

3.1 | Conveyor belts and luggage carousels

This challenge assumes a clear-cut distinction between memory and perception, with perception exclusively anchored in the fleeting present, and anything else confined to memory. The complaint is that the argument above assigns perception of objects through occlusion to the wrong side of that divide, because of too loose a temporal restriction on the causal relation in question. But is so strict a demarcation sustainable? Memory clearly supports perceptual representation, and perceptual representations are encoded in and may rely on memory. But how are we, more generally, to conceive of the relationship between them? To what extent are these categories opposed to one another, and to what extent are they overlapping or mutually constitutive?

On a flat-footed picture of the relationship between perception and memory there is a conveyer belt from one to the other. Perception involves the representation of a stimulus that is occurrently reflecting light into the eye. On this model, those representations are "briefly maintained in a sensory register…., processed in short-lived memory stores" (including transsaccadic memory, conceptual short-term memory and visual short-term memory), and then "[i]f attention is maintained… consolidated in long-term memory" (Intraub & Dickinson, 2008, 1012).[23] On this approach, what you can be said to *see* at any given moment is what is currently reflecting light into the eye *at that moment*, or at least within the few hundred milliseconds that it takes you to perform the neural processing required to generate a visual representation on the basis of those signals. Everything else is a form of memory.

The conveyor belt model has epistemic implications: as representations are shunted out of the sensory workspace into increasingly stable forms of memory storage, they become vulnerable to increasing levels of corruption, extra-sensory influences which impair the accuracy of the resulting representation. Its more intimate connection to the stimulus is what gives epistemic privilege to the perceptual representation, because *you can't have perception without the stimulus*, whereas you can have memories in the absence of the stimulus itself.

The conveyor belt picture, and the objection it supports, implies a simple contrast between memory and perception: a representation is *either* perceptual, or stored in memory. Representations move from one to another. But our growing knowledge of the processes that underlie visual perception puts pressure on that simple distinction: many visual experiences *rely* on contents stored in memory, and at small temporal scales there is often no easy way of demarcating working memory from perception.

It is well-established, for instance, that the visual system draws on stored information in the form of priors, to resolve ambiguity in incoming retinal information. Those expectations about the environment can be modified on the basis of previously encountered information. For instance, the ambiguity between distance and size is resolved in part on the basis of stored information about the kinds of objects we expect to encounter in a given environment, and our familiarity with their standard size (Davenport & Potter, 2004; Palmer, 1975). These priors encode those previously encountered regularities. In this way, they are a form of long-term memory.[24]

Relatedly, there is extensive evidence that information stored in both working and longer-term memory influences the content of conscious perceptual experience. This is most easily detected when it gives rise to perceptual illusion: in color-memory effects, for instance, stored information about the normal color of bananas and other objects influences subjects' perception of the color of images of them (Hansen, Olkkonen, Walter, & Gegenfurtner, 2006).[25]

This kind of to-and-fro of information between memory and perception has led to ambiguity over whether certain illusions are properly seen as features of *memory* or of perception itself. 'Boundary extension' refers to the way in which subjects falsely report having seen beyond the edges of an image in a picture, or the limits of their own visual field. If an image includes the middle portion of a chair and table, say, viewers typically report having also seen the extremities of those objects. Intraub and Dickinson (2008) found that this effect arises extremely rapidly – as soon as $1/20^{th}$ of a second after exposure to an image. So fast is the construction of the illusory content that they propose that it may instead be part of the initial perceptual representation of the scene, and that that perceptual representation is not derived from a single source of sensory input, but that it simultaneously draws on internally generated stores of information.

Boundary extension also helps us understand the ecological role of that blending between memory and perception: you can only look at a very small part of your environment at a time. Constructing a representation of the periphery of your view, in a manner that integrates sensory signal with information stored from moments before about how those now-invisible parts of the scene look, helps you to act successfully in a 360^{o} world. For our purposes, the final designation of the phenomenon is not as important as the possibility of dispute, and the support it gives the claim that working memory plays an important and routine role in determining what you see.

In light of these sorts of effects, rather than a conveyor belt, the relationship between perception and memory looks more like a luggage carousel: instead of a linear progression of information from sensory registration through increasingly stable forms of memory storage, what we see in perception is the repeated circling of some pieces of information whilst others come and go. Items are added to the carousel as new sensory information is registered, and representations leave the carousel as they cease to be maintained in a sensory register. But within the category we intuitively recognize as perceptual, experience relies in a wide range of ways on information stored in memory.

3.2 | Occurrent retinal stimulation as a criterion of seeing

These effects show that working memory can *blend with* sensory information to produce a conscious experience. But in the case of invisibility, the subject's experience of the relevant object would need to be *entirely* supported by memory at a moment in time. And that seems to flout an obvious way of distinguishing visual perception and memory: the former cannot arise in the absence of retinal stimulation by light reflected from the object, whereas memory can. Hence, whatever object representations can be maintained in the absence of such stimulation are mnemonic rather than visual.

Unfortunately, the case of saccades undermines the viability of occurrent dependence on retinal stimulation as a criterion of seeing. When saccading you rely on memory to integrate distinct frames of fixation: memory of what you saw in the moments immediately before lets you pick up where you left of, making sense of the new information by integrating it with the old into a coherent, detailed model of your surroundings (Hollingworth et al., 2008). Moreover, during the brief period of saccadic suppression, visual persistence (the maintenance of a visual image that persists after the cessation of light signals on the retina), is achieved by "filling in" the experience with information gathered a moment before, allowing us not to notice the continual interruption of our experience by these events in which the sensory signal is interrupted.

During a saccade, your experience is briefly independent of occurrent sensory stimulation: it relies on information from a moment before. A view that claims such stimulation is essential for perception is forced to say that our conscious experience flips constantly between moments of seeing and moments of remembering. And yet to the subject these states are functionally and phenomenally indistinguishable. The occurrent stimulus criterion leaves us then with a boundary between perception and memory that fails to track a distinction of theoretical or functional significance. It would be highly artificial to insist on segmenting our ordinary experience into milliseconds of seeing and milliseconds of remembering. Rather than insisting that you *stop seeing* whilst you are remembering, we should allow that memory supports a temporally extended visual experience that continually blends moments of occurrent sensory stimulation with stored information from moments before.

Can we scale this claim up to the case of occlusion, where the subject is *consciously aware* of the momentary invisibility of the object? In one sense, the subject's representation is clearly visual, since they represent the object *as occluded* and the relationship between subject and occluder is not subject to any temporal delay. But this might be taken to support merely a visual representation of the occluder, not of the object itself, on the grounds that that the period of independence of the *object representation* from retinal stimulation is too long. In response to this, the case of saccades suggests that temporal dislocation *per se* does not provide a criterion that can successfully distinguish memory and perception. The question then is what *degree* of temporal dislocation is tolerable. There is unlikely to be a clear cut off. But a dynamic priority claim licenses a greater degree of flexibility in the causal relation between object and experience: we plausibly need causal contact *at some point* during the extended period of time that has priority in determining the content and character of the moment in question. That is satisfied in the case of the tunnel effect: the object provides proximate retinal stimulation prior to its momentary occlusion, which continues to influence the content and character of the experience during its brief period of invisibility, as reflected in the experience of it as passing behind the occluder.

3.3 | A neural criterion?

Could facts about neural activity provide a simple way of teasing apart visual experiences from memories? It is natural to hope that mapping which areas of the brain are particularly active whilst the subject sustains a given representation could provide a clearer criterion, and one that might give grounds to dismiss the representation of occluded objects as an instance of memory rather than perception.

Appealing to a neural criterion relies on a prior theoretical commitment, to allowing facts about brain activity to guide our psychological distinctions, but that is a commonplace enough assumption. We know, for instance, that activity in the visual cortex, areas V1 through V4, supports visual perception. The hippocampus plays a particular role in serving explicit episodic and semantic memory, with implicit memory relying on the amygdala, and the prefrontal cortex subserving working memory specifically. If the maintenance of representations during occlusion relied not on the visual cortex but on areas more commonly associated with memory, that might give us some reason to regard extended object representations during occlusion as a form of memory rather than perception.

The first difficulty with this approach is that the crucial underlying assumption on which it rests, namely that there is a clear or consistent distinction between the areas of the brain responsible for visual representation and representation in memory, is false. There is good evidence that neural activity in the visual cortex is, on occasion, common to the maintenance of representations held in working memory over periods of several seconds in the absence of a physical stimulus, as well as to visual perception. Take, as an example of such work, a study by Harrison and Tong (2009). Harrison and Tong showed participants two gratings that varied in their orientation, one after the other. They then provided them with a numerical cue, indicating that they should remember either the first or the second of the gratings they had been showed. After an 11 second retention interval, a test grating was presented, and participants were tasked with indicating which way it was rotated, relative to the cued grating. They found that observers performed equally well, regardless of which grating was cued, and that they were capable of discriminating small differences in orientation between the cued grating and the test, indicating that they were successfully remembering the stimuli. But most significantly, they found that "ensemble activity pooled from areas V1-V4, was highly predictive of the orientation held in working memory, with prediction accuracy reaching 83%" (2009, 633). Activity in those areas is normally strongly associated with visual perception. The fact that pooled information from those areas allowed the experimenters to predict which orientation was held in working memory suggests that those neural populations taken to underwrite visual representations play an important role in supporting representations in memory too. Their conclusion was that "retaining an orientation in working memory recruits many of the same orientation-selective subpopulations as those that are activated under stimulus-drive conditions"(2009, 633). The authors take their results to show that "early visual areas are not only important for processing information about the immediate sensory environment, but can also maintain information in the absence of direct input to support higher-order cognitive function" (2009, 634).[26] If that is right, then there is no clear or abrupt distinction even in the neuronal populations that support perception and memory.

But this first difficulty pales when placed beside a second: the interpretation of these kinds of results itself relies on substantive theoretical commitments. Take work by Hulme and Zeki (2007) looking at neural activity in areas of the brain associated with the perception of particular objects during their occlusion. Perception of faces often correlates with activity in the fusiform

face area, for instance, whilst the lateral occipital complex is known to play a particular role in the visual perception and recognition of objects. Hulme and Zeki showed subjects displays in which an image of a face or a house either remained visible whilst a screen rose behind it, or was occluded to 98% by the screen rising in front of it. They then measured activity in brain regions over the next 7.5 seconds whilst subjects continued to look at the occluded object. They found that "activity within the FFA and LOC is invariant to whether objects are occluded or not.... The surprise lies in the fact that the areas were activated with the same magnitude and had very similar time courses, whether perceived or not" (2007, 1201).

On the basis of this, Hulme and Zeki conclude that activity in these regions indicates *awareness of presence* rather than perception, as had previously been assumed. But an alternative interpretation is available: that activity in these areas does indeed indicate perception, and that the objects in question continue to be perceived even as most of their surface is not currently visible. In these cases, subjects had no occurrent stimulation. Whatever representations subjects maintained of the objects in question relied on memory, but it was nonetheless reflected in activity in regions of the brain traditionally associated with perception.

Might Hulme and Zeki's argument implicitly appeal to the following more sophisticated form of neural criterion: genuine perception requires not just activation of the visual cortex, but activation *throughout* the standard set of neural areas associate with perception, that is, the retina, the optic nerve, the magno- and parvocellular pathways in addition to the visual cortex. What makes representations of invisible objects instances of memory rather than perception, as during occlusion or saccadic suppression, is that there is a lack of activity *at some point* in that chain, even if certain parts of the visual cortex continues to be recruited in service of these representations.[27]

Setting aside the difficulty of arriving at a non-question begging description of those areas, this proposal ultimately serves to highlight the limitations of the neural approach. In the first place, experience of occluded objects *is* consistent with stimulation of the full range of neuronal activity associate with perception, since the occluder continues to reflect light into the retina, even if the object behind it does not. And in the other direction, there seem to be moments of visual perception which fail to satisfy this criterion. Recall the case of visual experience during saccades. At these moments there is reduced activity in the magnocellular pathway. But we saw previously that it would be undesirable to classify the experience had during these moments as failing to constitute a moment of seeing. But if we count visual experience at moments of saccadic suppression as a form of seeing, then it looks like there is no clear paradigmatic case of perceptual experience that satisfies this proposed criterion. Even this more sophisticated neural criterion fails to successfully demarcate remembering from seeing.

3.4 | Taking stock

We have surveyed a range of ways in which one's visual experience at a moment is not, and cannot be, immediately derived from the stimulation of the retina at that moment. Long term memory is an essential ingredient, as priors resolve ambiguity in retinal input. Working memory routinely blends with sensory information to inform the content of visual experience. Transsacadic memory joins up brief spells of sensory information into a coherent experience. There is considerable overlap between the neural structures and activity associated with memory and perception. The upshot of all this for our purposes is that reliance on memory fails to provide a coherent reason for disqualifying occluded object apprehension from visual perception. Extended object perception relies on memory, but so do many other core aspects of perception.

4 | CONCLUSION

Conscious visual experience is a fusion of elements from short-term (working) memory, long-term memory (priors), and occurrent sensory stimulation. This thorough-going infusion of memory into perception makes it hard to draw a clean line between what you are currently seeing, and what you are merely remembering. Far from living in the present, an entwinement with the past is inherent in the most basic of perceptual capacities, our perception of objects in the world around us.

The involvement of memory provides no barrier to recognising extended object perception as a genuinely visual form of representation. Settling that clears the way for the claim that a subject sees the object as it passes behind the occluder, despite its momentary invisibility.

In some ways this is a highly surprising upshot. But it seems less counterintuitive if we relax a habitual insistence on the priority of the moment in visual perception. The object in question is only properly deemed invisible when we use the *moment* as the criterion for invisibility. If we instead prioritise a longer period of time, then the object *is* visible during that time frame. Read another way, the argument here encourages us to be dynamists about visibility, as well as about visual perception. When we do so, we treat visual perception as a modelling process that can at points operate independently of occurrent retinal input.

However we accommodate the claim that we see the object as it passes behind the occluder - whether we think of visibility as a dynamic property, or allow that we can see invisible objects - goes to the core of *what visual perception is*: accommodating the dynamism of visual experience forces us to think of it not just as a process of passive reflection such that any extended period of experience supervenes on its component instants, but as an extended process of active modelling of the environment on the basis of a particular set of light signals. That active modelling can continue, to a limited extent at least, during periods of occlusion, or other moments when less sensory information is available. Dynamic perception, properly understood, involves developing and maintaining models of the environment. Those models are largely driven by occurrent retinal stimulation, but that retinal stimulation is used to best purpose when it forms the basis of more extensive representations capable of integrating information gathered at different moments in time.

It is commonplace to emphasise the challenge of visually representing a 3D world on the basis of 2D light signals. That underdetermination is certainly part of what mandates a reliance on stored information. Less emphasis has been given to another challenge the visual system faces: that of representing a dynamic, changing environment that extends to 360 degrees whilst only being able to direct its gaze towards one direction at a time. To solve that more comprehensive problem, the visual system must be a skillful quilter, uniting spatially and temporally disparate pieces of retinal information into a representation that is coherent and unified across both time and space.[28]

NOTES

[1] Note that we are not (primarily) interested in the question of whether visual perceptual *processing* is temporally extended. Visual processing in creatures like ourselves takes time. That does not decide the question of dynamic priority. It is possible that the result of such a temporally extended process could be a snapshot with no duration itself, just as a still life painting that captures the rotting fruit bowl at a single moment on a heavy summer afternoon may have taken a much longer time to produce.

[2] See Chuard (2011, 15) for an articulation of dynamic priority as the claim that the experience fails to supervene on a succession of its parts.

[3] It is tempting to confuse the question of whether the most basic unit of visual experience is temporally extended with the question of whether visual experience is discrete or continuous, but the two questions come apart. Whilst it can seem natural to suppose that snapshots would be discrete, if continuity is understood merely the *absence of gaps*, and if experience is potentially composed of an *infinite* sequence of snapshots, then visual experience could be a continuous sequence of snapshots. Equally, the dynamist can also allow that there are privileged joints in the flow of visual experience that allow us to distinguish extended experiential atoms from one another. See Lee (2014) whose retentionalism is by his own description a form of "extended atomism".

[4] See Dainton, 2000 for a much fuller exposition. Some extensionalists, like Phillips, in addition endorse an inheritance principle: not only are temporal experiences themselves temporally extended, but "for any temporal property apparently presented in perceptual experience, experience itself has that same temporal property" (Phillips, 2014, 131). An experience of one event succeeding another, for instance, entails that the experience of the second event succeeds the experience of the first (though Phillips also denies that temporal content is *explained* by the temporal properties of the experience).

[5] I also take Dynamic Priority to be incompatible with a variant of the snapshot or cinematic model, the so-called *dynamic* snapshot view endorsed by Prosser (2016) and Le Poidevin (2007) among others, according to which snapshots can include dynamic content, without the resulting experience having any temporal breadth. See Shardlow (2019) for an exposition and critique of these views.

[6] Again, questions about the extension of the subject's experience should be distinguished from the period of time over which the perceptual system gathers the information which it integrates in order to arrive at a subjective experience – the window of temporal integration (Forget, Buiatti, & Dehaene, 2009; van Wassenhove, Grant, & Poeppel, 2007). The snapshot view of experience is quite compatible with a view on which those snapshots are based on an extended window of temporal integration.

[7] Whitney Houston and Beyonce's renditions of the Stars and Stripes at the 1991 and 2016 Superbowls are good examples of a perfect and an interrupted cadence respectively.

[8] See Tyler Burge (2010, 450–65) for the claim that a capacity to perceive objects as temporally extended is a necessary prerequisite for object perception.

[9] We visually experience objects in the environment as having a constant shape and colour, though the proximal retinal stimulation they give rise to changes moment by moment as we and they move, and light levels change. The phenomenal primacy of perceptual constancy is at least suggestive of a dynamic priority claim, as it can only emerge over a protracted period of time. See Hatfield and Allred (2012) for further empirical discussion of perceptual constancy.

[10] More controversially, objects may be invisible in this sense when they are transparent, or when it is dark. For an argument that we see objects in the dark that provides an alternative route to *Seeing the Invisible* see Sorensen (2004).

[11] The extent of saccadic suppression depends on spatial frequency: high spatial frequencies may become invisible whilst lower spatial frequencies become more conspicuous, suggesting that saccadic suppression may be specific to the magnocellular pathway, and effect motion sensitivity in particular (Ross, Morrone, Goldberg, & Burr, 2001).

[12] Is that diminution in sensitivity to location inconsistent with reasonable sensitivity to the object's appearance? There is good reason to think not. Attentional effects introduce widespread though minor distortions in spatial perception in ways that have implications for the representation of object location (Gobell and Carasco 2005, Liverence & Scholl, 2011). Imposing too high a standard of sensitivity to object location risks excluding many plausible cases from satisfying the criterion of seeing. I say more about what the requirement of reasonable sensitive entails below in section 2.4.

[13] I am indebted to an anonymous reviewer for suggesting a similar case to this.

[14] Interested readers can compare the relevant cases here: http://perception.yale.edu/Brian/demos/Tunnel-CD.html

[15] For how long do object files persist? Noles, Scholl, and Mitroff (2005) found evidence that they can persist for at least 8 seconds. They write that "[o]bject files may be the "glue" that makes visual experience coherent not just in online moment-by-moment processing, but on the scale of seconds that characterizes our everyday perceptual experiences" (2005, 324).

[16] Does appealing to object files to support a claim about the nature of perceptual representation involve a vehicle/content confusion? Dennett and Kinsbourne warn that "In general, we must distinguish features of

representings from the features of representeds; someone can shout "softly, on tiptoe" at the top of his lungs, there are gigantic pictures of microscopic objects, and oil paintings of artists making charcoal sketches. The top sentence of a written description of a standing man need not describe his head, nor the bottom sentence his feet. To suppose otherwise is to confusedly superimpose two different spaces: the representing space and the represented space"(1992, 188). Though their point is well taken, it's equally noticeable that the vehicle of representation *does* constrain the content in various ways. The content of an oil painting may be limited by its size and shape. A collection of several oil paintings make different kinds of representation possible than just a single painting can on its own. And there are certain forms of content which are entirely precluded by the medium of an oil painting. For this reason, it can be legitimate to look to features of the vehicle to tell us something about the parameters within which visual representation operates. The temporal persistence of object files, our best psychological model of the underlying mechanism of object representation, is an important mechanism by which information gathered over an extended period of time can be integrated to inform the contents of a temporally extended conscious experience. The presence of that mechanism at least *allows* for greater flexibility between the time at which retinal information is registered and its manifestation in a conscious visual experience.

[17] There is evidence of a capacity limit of roughly 4 on the number of objects that we can successfully track at a time, possibly reflecting the fact that maintaining these files taxes limited attentive resources (Scholl and Pylyshyn 1999, though see Alvarez & Franconeri, 2007 for evidence of greater flexibility in these limits). We frequently attend to some objects whilst continuing to see others in the background. In such a case we are unlikely to maintain object files for those unattended objects given the limit on how many such files we are capable of simultaneously maintaining. This is evidenced by our difficulties detecting featural changes in cases of inattentive change blindness (Simons and Levin 1997).

[18] This invites the question: where *should* we set the relevant bar? How much sensitivity to the object's appearance does the criterion of seeing require? It is unlikely that there will be any clear cut off here, particularly since the appropriate vividness of the representation of the object will depend on contextual factors including distance and climate, for instance. The relevant standard may in part be a negative one: reasonable sensitivity is inconsistent with *gross mis*representation of an object's features though compatible with agnosticism about them, as when the subject is seen at a distance, or with errors about several individual properties.

[19] Lande (ms.) quotes Child (1992, 311) voicing the claim that this is in fact a condition on conceptual possession: "If one has the concept of vision, one knows that S will stop seeing something if he shuts his eyes, or if we interpose something opaque between him and the object."

[20] See Lande (ms.) for interesting discussion of relevant empirical work by Gold, Murray, Bennett, and Sekuler (2000), that suggests subjects viusally interpolate occluded countours of an object, in a way that can be disrupted by visual noise, despite the invisibility of the relevant parts of the objects.

[21] Note, though, that the subject who *sees* the object behind the occluder *is* accurately representing the world before them: they represent the fact the object is *still there* but that it is temporarily occluded by another object.

[22] See Lande (ms.) for an argument that visual reference extends beyond seeing to include cases where (parts of) the target objects are not visible.

[23] Intraub and Dickinson offer this articulate description of a traditional picture which they themselves challenge.

[24] Even the prior that light comes from above has been shown to be modifiable by interactive experience within an environment (Adams, Graf, & Ernst, 2004)

[25] See also Mendoza, Schneiderman, Kaul, and Martinez-Trujillo (2011) and Kang, Hong, Blake, & Woodman, 2011) for work that purports to show the influence of contents stored in short-term working memory on conscious visual experience.

[26] Could we test the claim that visual perception is thoroughgoingly reliant on memory, by looking to see if amnesia is associated with deficits in visual perception? Unfortunately, the answer to this is complex: amnesia *is* often associated with deficits in perception, but it is unclear whether that is because memory subserves perception, or because the kinds of neural deficits responsible for the amnesia are also responsible for independent impairments to perception. For instance, Lee et al. (2005) found that damage to the hippocampus, an area of the brain known to underwrite various forms of memory, was accompanied by poorer skills at visually discriminating spatial scenes. That finding could be interpreted either as evidence for the role of memory in perceiving spatial scenes, *or* as evidence of the role of the hippocampus in supporting spatial perception independently of the role it plays in memory.

[27] I am grateful to an anonymous reviewer for directing my attention to this possible criterion.

[28] For discussion of earlier drafts of this paper I am grateful to audiences at CUNY, the 2018 meeting of the Society for Philosophy and Psychology, Oxford's Jowett society, a workshop on perception at the University of Glasgow, and the University of Texas at Austin MindWork seminar. Thanks to Greyson Abid, Richard Holton, Ian Phillips, Jake Quilty-Dunn, Jack Shardlow, Susanna Siegel, and Kate Vredenburgh for comments on or discussion of various parts of this material. I am particularly indebted to an anonymous reviewer at this journal for their extensive comments which greatly improved the paper.

REFERENCES

Adams, Wendy J., Graf, Erich W., & Ernst, Marc.O. (2004). Experience Can Change the 'Light-from-above' Prior. *Nature Neuroscience*, *7*, 1057–1058.

Alvarez, George, & Franconeri, Steven (2007). How many objects can you track?: Evidence for a resource-limited attentive tracking mechanism" *Journal of Vision*, *7*(13). 14, 1–10 https://doi-org.ezp.lib.cam.ac.uk/10.1167/7.13.14

Burge, Tyler. (2010). *Origins of Objectivity*. Oxford University Press. Retrieved from https://www.oxfordscholarship.com/view/10.1093/acprof:oSo/9780199581405.001.0001/acprof-9780199581405

Burke, Luke. (1952). On the Tunnel Effect. *Quarterly Journal of Experimental Psychology*, *4*(3), 121–138. https://doi.org/10.1080/17470215208416611

Child, William. (1992). Vision and Experience: The Causal Theory and the Disjunctive Conception. *The Philosophical Quarterly*, *1950-) 42*(168), 297–316. https://doi.org/10.2307/2219682

Chuard, Philippe. (2011). Temporal Experiences and Their Parts. *Philosopher's Imprint*, *11*(11). Retrieved from http://hdl.handle.net/2027/spo.3521354.0011.011

Colman, Andrew M. (2009). "Apparent Movement." In *A Dictionary of Psychology*. Oxford University Press. Retrieved from https://www.oxfordreference.com/view/10.1093/acref/9780199534067.001.0001/acref-9780199534067-e-580

Dainton, Barry. (2000). *Stream of Consciousness: Unity and Continuity in Conscious Experience*. Routledge.

Dainton, Barry. (2018). "Temporal Consciousness." In *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (Ed.), Winter 2018. Metaphysics Research Lab, Stanford University. Retrieved from https://plato.stanford.edu/archives/win2018/entries/consciousness-temporal/

Davenport, J. L., & Potter, M. C. (2004). Scene Consistency in Object and Background Perception. *Psychological Science*, *15*, 559–564.

Dennett, Daniel C., & Kinsbourne, Marcel (1992). Time and the Observer: The Where and When of Consciousness in the Brain. *Behavioral and Brain Sciences*, *15*(2), 183–247. https://doi.org/10.1017/S0140525/00068229.

Flombaum, Jonathan I., & Scholl, Brian J. (2006). A Temporal Same-Object Advantage in the Tunnel Effect: Facilitated Change Detection for Persisting Objects. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(4), 840.

Forget, Joachim, Buiatti, Marco, & Dehaene, Stanislas (2009). Temporal Integration in Visual Word Recognition. *Journal of Cognitive Neuroscience*, *22*(5), 1054–1068. https://doi.org/10.1162/jocn.2009.21300

Gobell, Joetta, & Carrasco, Marisa (2005). Attention Alters the Appearance of Spatial Frequence and Gap Size" *Psychological Science*, *16*(8), 644–651 https://doi-org.ezp.lib.cam.ac.uk/10.1111/j.1467-9280.2005.01588.x

Gold, Jason M., Murray, Richard F., Bennett, Patrick J., & Sekuler, Allison B. (2000). Deriving Behavioural Receptive Fields for Visually Completed Contours. *Current Biology*, *10*(11), 663–666. https://doi.org/10.1016/S0960-9822(00)00523-6.

Green, Edwin James, & Quilty-Dunn, Jake (2017). What Is an Object File?" *The British Journal for the Philosophy of Science*,

Greenlee, Mark W., & Kimmig, Hubert (2019). "Visual Perception and Eye Movements." In *Eye Movement Research: An Introduction to Its Scientific Foundations and Applications*. Christoph Klein & Ulrich Ettinger (Eds.), (pp. 165–196). Studies in Neuroscience, Psychology and Behavioral Economics. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-20085-5_5

Grush, Rick. (2005). "Brain Time and Phenomenological Time." In *Cognition and the Brain: The Philosophy and Neuroscience Movement* (Vol. 160). Andrew Brook & Kathleen Akins (Eds.), Cambridge University Press.

Hansen, Thorsten, Olkkonen, Maria, Walter, Sebastian, & Gegenfurtner, Karl R. (2006). Memory Modulates Color Appearance. *Nature Neuroscience*, *9*(11), 1367–1368. https://doi.org/10.1038/nn1794

Harrison, Stephenie A., & Tong, Frank (2009). Decoding Reveals the Contents of Visual Working Memory in Early Visual Areas. *Nature*, *458*(7238), 632–635. https://doi.org/10.1038/nature07832

Hatfield, Gary, & Allred Sarah eds., (2012). *Visual Experience: Sensation, Cognition, and Constancy*. Oxford, New York: Oxford University Press.

Hollingworth, Andrew, Richard, Ashleigh M., & Luck, Steven J. (2008). Understanding the Function of Visual Short-Term Memory: Transsaccadic Memory, Object Correspondence, and Gaze Correction. *Journal of Experimental Psychology. General*, *137*(1), 163–181. https://doi.org/10.1037/0096-3445.137.1.163

Hulme, Oliver J., & Zeki, S. (2007). The Sightless View: Neural Correlates of Occluded Objects. *Cerebral Cortex*, *17*(5), 1197–1205. https://doi.org/10.1093/cercor/bhl031

Intraub, Helene, & Dickinson, Christopher A. (2008). False Memory 1/20th of a Second Later: What the Early Onset of Boundary Extension Reveals About Perception. *Psychological Science*, *19*(10), 1007–1014. https://doi.org/10.1111/j.1467-9280.2008.02192.x

Kahneman, Daniel, Treisman, Anne, & Gibbs, Brian J. (1992). The Reviewing of Object Files: Object-Specific Integration of Information. *Cognitive Psychology*, *24*(2), 175–219. https://doi.org/10.1016/0010-0285(92)90007-O

Kang, Min-Suk, Hong, Sang Wook, Blake, Randolph, & Woodman, Geoffrey F. (2011). Visual Working Memory Contaminates Perception. *Psychonomic Bulletin & Review*, *18*(5), 860–869. https://doi.org/10.3758/s13423-011-0126-5

Krekelberg, Bart. (2010). Saccadic Suppression. *Current Biology*, *20*(5), R228–R229.

Lande, Kevin. ms "Seeing and Visual Reference."

Le Poidevin, Robin. (2007). The Images of Time: An Essay on Temporal Representation. The Images of Time. Oxford University Press. Retrieved from https://www.oxfordscholarship.com/view/10.1093/acprof:oSo/9780199265893.001.0001/acprof-9780199265893

Lee, Andy C. H., Bussey, Tim J., Murray, Elisabeth A., Saksida, Lisa M., Epstein, Russell A., Kapur, Narinder, … Graham, Kim S. (2005). Perceptual Deficits in Amnesia: Challenging the Medial Temporal Lobe 'Mnemonic' View. *Neuropsychologia*, *43*(1), 1–11. https://doi.org/10.1016/j.neuropsychologia.2004.07.017

Lee, Geoffrey. (2014). Temporal Experience and the Temporal Structure of Experience. *Philosopher's Imprint*, *14*(3). Retrieved from http://hdl.handle.net/2027/spo.3521354.0014.003

Li, Wenxun, & Matin, Leonard (1990). The Influence of Saccade Length on the Saccadic Suppression of Displacement Detection. *Perception & Psychophysics*, *48*(5), 453–458. https://doi.org/10.3758/BF03211589

Liverence, Brandon, & Scholl, Brian (2011). Selective Attention Warps Spatial Reservation: Parallel but Opposing Effects on Attended Versus Inhibited Objects. *Psychological Science*, *22*(12), 1600–1608 https://doi-org.ezp.lib.cam.ac.uk/10.1177/0956797611422543

Mendoza, Diego, Schneiderman, Megan, Kaul, Christian, & Martinez-Trujillo, Julio (2011). Combined Effects of Feature-Based Working Memory and Feature-Based Attention on the Perception of Visual Motion Direction. *Journal of Vision*, *11*(1), 11–11. https://doi.org/10.1167/11.1.11

Mitroff, Stephen R., Scholl, Brian J., & Wynn, Karen (2005). The Relationship between Object Files and Conscious Perception. *Cognition*, *96*(1), 67–92. https://doi.org/10.1016/j.cognition.2004.03.008

Noles, Nicholaus S., Scholl, Brian J., & Mitroff, Stephen R. (2005). The Persistence of Object File Representations. *Perception & Psychophysics*, *67*(2), 324–334. https://doi.org/10.3758/BF03206495

Nudds, Matthew. (2010). What Are Auditory Objects?" *Review of Philosophy and Psychology*, *1*(1), 105–122. https://doi.org/10.1007/s13164-009-0003-6

O'Callaghan, Casey. (2009). "Sounds." In *Oxford Companion to Consciousness*, Tim Bayne, Axel Cleeremans, & Patrick Wilken (Eds.), Oxford University Press. Retrieved from https://www.oxfordreference.com/view/10.1093/acref/9780198569510.001.0001/acref-9780198569510-e-308

Palmer, S. E. (1975). The Effects of Contextual Scenes on the Identification of Objects. *Memory and Congition*, *3*, 519–526.

Phillips, Ian B. (2010). Perceiving Temporal Properties. *European Journal of Philosophy*, *18*(2), 176–202. https://doi.org/10.1111/j.1468-0378.2008.00299.x

Phillips, Ian B. (2011). Perception and Iconic Memory: What Sperling Doesn't Show. *Mind and Language*, *26*(4), 381–411.

Phillips, Ian B. (2014). Experience of and in Time. *Philosophy Compass*, *9*(2), 131–144. https://doi.org/10.1111/phc3.12107

Prosser, Simon. (2016). Experiencing Time. Experiencing Time. Oxford University Press. Retrieved from https://www.oxfordscholarship.com/view/10.1093/acprof:oSo/9780198748946.001.0001/acprof-9780198748946

Reid, Thomas. (2015). In Derek R. Brooks & Knud Haakonssen (Eds.), *Essays on the Intellectual Powers of Man*. Edinburgh: Edinburgh University Press.

Ross, John, Morrone, M. Concetta, Goldberg, Michael E., & Burr, David C. (2001). Changes in Visual Perception at the Time of Saccades. *Trends in Neurosciences*, *24*(2), 113–121. https://doi.org/10.1016/S0166-2236(00)01685-4

Schellenberg, Susanna. (2016). Perceptual Particularity. *Philosophy and Phenomenological Research*, *93*(1), 25–54. https://doi.org/10.1111/phpr.12278

Scholl, Brian J. (2001). Objects and Attention: The State of the Art. *Cognition, Objects and Attention*, *80*(1), 1–46. https://doi.org/10.1016/S0010-0277(00)00152-9.

Scholl, Brian J., & Flombaum, Jonathan (2010). "Object Persistence." In *Encylopedia of Perception*, *2*, 653–657.Thousand Oaks, CA: Sage Publications.

Scholl, Brian J., & Pylyshyn, Zenon W. (1999). Tracking Multiple Items Through Occlusion: Clues to Visual Objecthood. *Cognitive Psychology*, *38*(2), 259–290. https://doi.org/10.1006/cogp.1998.0698

Scholl, Brian J., Pylyshyn, Zenon W., & Feldman, Jacob (2001). What Is a Visual Object? Evidence from Target Merging in Multiple Object Tracking. *Cognition*, *80*(1–2), 159–177. https://doi.org/10.1016/S0010-0277(00)00157-8.

Shardlow, Jack. (2019). Minima Sensibilia: Against the Dynamic Snapshot Model of Temporal Experience. *European Journal of Philosophy*, *27*(3), 741–757. https://doi.org/10.1111/ejop.12442

Simons, Daniel, & Levin, Daniel (1997). Change blindness. *Trends in Cognitive Science*, *1*.7:261–267

Sorensen, Roy. (2004). We See in the Dark. *Noûs*, *38*(3), 456–480. https://doi.org/10.1111/j.0029-4624.2004.00478.x

Spelke, Elizabeth S. (1988). "Where Perceiving Ends and Thinking Begins: The Apprehension of Objects in Infancy." In Perceptual Development in Infancy*., 197–234. The Minnesota Symposia on wo Child Psychology* (Vol. 20). Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc.

Vallines, Ignacio, & Greenlee, Mark W. (2006). Saccadic Suppression of Retinotopically Localized Blood Oxygen Level-Dependent Responses in Human Primary Visual Area V1. *The Journal of Neuroscience*, *26*(22), 5965–5969. https://doi.org/10.1523/JNEUROSCI.0817-06.2006

VanRullen, Rufin, & Thorpe, Simon J. (2001). The Time Course of Visual Processing: From Early Perception to Decision-Making. *Journal of Cognitive Neuroscience*, *13*(4), 454–461. https://doi.org/10.1162/08989290152001880

Vickery, Timothy J., & Chun, Marvin M. (2010). Object-Based Warping: An Illusory Distortion of Space Within Objects. *Psychological Science*, *21*(12), 1759–1764. https://doi.org/10.1177/0956797610388046

Wassenhove, Virginie van, Grant, Ken W., & Poeppel, David (2007). Temporal Window of Integration in Auditory-Visual Speech Perception. *Neuropsychologia, Advances in Multisensory Processes*, *45*(3), 598–607. https://doi.org/10.1016/j.neuropsychologia.2006.01.001

How to cite this article: Munton J. How to see invisible objects. *Noûs*. 2022;*56*:343–365. https://doi.org/10.1111/nous.12360