

UNCONSCIOUS PERCEPTION AND PHENOMENAL COHERENCE

Jake Quilty-Dunn
University of Oxford

Forthcoming in *Analysis*

Abstract: It is an orthodoxy in cognitive science that perception can occur unconsciously. Recently, Hakwan Lau, Megan Peters, and Ian Phillips have argued that this orthodoxy may be mistaken. They argue that many purported cases of unconscious perception fail to rule out low degrees of conscious awareness while others fail to establish genuine perception. This paper presents a case of unconscious perception that avoids these problems. It also advances a general principle of “phenomenal coherence” that can insulate some forms of evidence for unconscious perception from the methodological critiques of Lau, Peters, and Phillips.

§1. Introduction

Long thought to be a contradiction in terms, the existence of unconscious perception became orthodoxy in cognitive science in the latter half of the twentieth century. Recently, theorists such as Hakwan Lau, Megan Peters and Ian Phillips have argued that this orthodoxy may be mistaken.

Evidence for unconscious perception typically involves above-chance performance on a task that requires sensitivity to some stimulus accompanied by failure to report perceiving the stimulus (e.g., Marshall & Halligan 1988; Luck et al. 1996). The recent bevy of critics have raised two problems with this form of evidence. The first is known as the ‘criterion problem’—namely, that failures of report may be a result of an overly conservative criterion for report rather than a lack of consciousness. Peters and Lau, for instance, write that ‘an observer’s report of “unseen” doesn’t necessarily imply *complete* lack of awareness, only that the stimulus’ strength fell below some arbitrary boundary for reporting “seen” (2015: 1). A subject might be dimly aware of the stimulus but not feel sure enough to report it.

This problem threatens the orthodoxy since, if it indeed highlights a methodological flaw, then it cuts down large swaths of evidence for unconscious perception in one fell swoop. If Phillips is right that ‘we need to consider seriously the possibility that the response pattern reflects conscious perception unreported due to a conservative response criterion’ (Phillips 2016: 428), then there may not be any evidence for unconscious perception after all (cf. Block 2016). And if there is no unconscious perception, then we need to rethink the most basic aspects of perception and its relation to the rest of the mind. For example, theories of consciousness are typically formulated in terms of differences between conscious and unconscious perception (Prinz 2015), and unconscious perception may support representational over relational views of perception (Berger & Nanay 2016; cf. Anaya & Clarke 2017). If perception is intrinsically conscious, then much philosophical work on the nature of perception and consciousness rests on a mistake.

Phillips and others have also raised the problem that demonstrating unconscious sensory processing does not necessarily demonstrate unconscious perception—what Phillips calls the ‘problem of attribution’ (Phillips forthcoming). Virtually all theorists would grant that transduction is unconscious and that at least some processing within the visual system occurs unconsciously (Phillips 2016: 419). But mere sensory processing of this sort is arguably insufficient for genuine perception (cf. Prinz 2015). Genuine perception is personal-level rather than subpersonal—i.e., attributable to the subject rather than merely to her visual system.

These two problems come apart from one another. An experiment may avoid the criterion problem, thereby demonstrating unconscious processing, but fail to show that this processing constitutes genuine perception. Another may demonstrate genuine personal-level perception, and thus avoid the problem of attribution, but fail to show that the lack of report is due to unconsciousness rather than a conservative report criterion. The essence of the critique from Phillips, Lau and Peters is that no putative evidence for unconscious perception successfully overcomes both problems at once.

The proximal goal of this paper is to provide a specific example of unconscious perception that avoids these pitfalls. The distal goal is to push the debate forward by sketching a methodological approach that accommodates the insights of the aforementioned critics and may provide further evidence for the existence of unconscious perception.

§2. Object files and phenomenal coherence

The illustrative example I will use is a study by Stephen Mitroff, Brian Scholl and Karen Wynn (2005). Mitroff et al. used an ‘object-reviewing’ paradigm. In typical object-reviewing paradigms, some small number of stationary objects (usually two) is displayed, and a feature (e.g., a letter) is displayed within the boundaries of each object and thus ‘previewed’ (Kahneman et al. 1992). Once the previewed features disappear, the objects move to different locations. A feature then appears in one object and participants indicate whether this test feature matches one of the previewed features (in some variants, participants simply name the test feature).

What has been found and replicated many times over is an *object-specific preview benefit* (OSPB)—that is, participants more quickly identify a match when the test feature appears in the same object in which it was previewed. Since the objects are qualitatively identical and move locations, this preview benefit is object-specific rather than feature-specific or location-specific (though there is location-specific priming as well [Hollingworth & Rasmussen 2010]). The visual system deploys a perceptual object representation or *object file* that represents each object, stores information about its features and tracks it as it moves (Green & Quilty-Dunn forthcoming; Quilty-Dunn ms).

Mitroff et al. (2005) ran a normal object-reviewing experiment with one quirk: the objects’ motion was visually ambiguous between ‘bouncing’ off each other, or ‘streaming’ through each other (Fig. 1). Participants completed an object-reviewing task and indicated whether they consciously experienced bouncing or streaming. Remarkably, while subjects indicated a conscious experience of streaming in virtually all trials (e.g., 97% in one experiment), the OSPB showed that the object files tracked the objects as though they had bounced. The authors ran follow-up experiments controlling for variables like distance and motion direction, thereby eliminating alternative explanations and replicating the effect.




Figure 1 — From Mitroff et al. 2005

The evidence strongly suggests that the object files in these experiments are unconscious. The object files represented objects as bouncing off each other, which directly contradicts the experience of objects as streaming through each other. This aspect of the experiment, I argue, enables it to avoid the criterion problem.

In experiments targeted by the criterion problem, subjects fail to report that they see a stimulus or notice a difference between two stimuli. For example, Phillips (2016) illustrates the criterion problem by appeal to the famous case study from Marshall and Halligan (1988). They presented a subject with hemispatial neglect (i.e., an apparent inability consciously to perceive stimuli on one side of the visual field) with two pictures of houses. One house was engulfed in flames, but only on the neglected side. The patient reported seeing no difference between the houses, yet when asked in which house she would prefer to live, she consistently chose the house that was not on fire. Phillips argues that the inference that the patient lacked conscious perception of the flames ignores the possibility ‘that she has degraded conscious perception but that such perception falls below her response threshold’ (2016: 432–33).

The criterion problem arises in circumstances where subjects fail to indicate that they have seen a stimulus though they may have had dim conscious awareness of it. Peters, for example, writes in a discussion of the criterion problem that ‘for a stimulus to be unconscious, an observer’s subjective experience of the stimulus should be no different from the subjective experience of nothing at all’ (Peters et al. forthcoming; cf. Block in Peters et al. forthcoming). The criterion problem plagues experiments that fail to distinguish dim perceptual phenomenology from the absence of perceptual phenomenology.

The Mitroff et al. experiment, on the other hand, does not fall prey to this problem. There is no concern that subjects fail to indicate how they see the objects moving—on the contrary, the subjects indicated that they saw the objects stream through each other on as many as 97% of trials. Unlike in Marshall and Halligan’s (1988) study and other studies critiqued by Phillips, Lau and Peters, the issue is not whether the subjects are willing to report detection of the stimulus. They do indicate conscious detection of the stimulus, and the experience they describe is plainly at odds with their object files; thus, the object files do not occur consciously.

There is a reasonable methodological principle to be invoked here, which I shall call the Principle of Phenomenal Coherence (PPC):

If (i) a subject visually experiences F and (ii) also visually represents G and (iii) F and G are clearly incompatible, then, *ceteris paribus*, there is good reason to think that the representation of G is unconscious.

Consider cases of binocular rivalry, wherein different stimuli are presented to each eye (e.g., a face and a house). Rather than causing an incoherent conscious experience, phenomenology flips back and forth between coherent experiences of each stimulus while the other stimulus is merely registered unconsciously (Blake & Logothetis 2002). Cases of this sort suggest that conscious visual experience typically involves the univocal selection of a coherent percept rather than conscious co-occurrence of contradictory contents.

Endorsing PPC does not require denying that conscious experiences ever have contradictory contents. There are purported cases of contradictory contents in experience, such as the waterfall illusion (Crane 1988). If you take a long look at a continually moving stimulus and then look at a stationary object, perceptual adaptation will (so the conventional wisdom goes) cause you illusorily to see the object as moving while simultaneously seeing it as stationary relative to its surroundings. Moreover, both these contents are consciously experienced.

A contradictory experience like the waterfall illusion, however, is jarring, even disorienting, in its manifest incoherence. The *ceteris paribus* clause of PPC should therefore contain the provision that the relevant experience does not exhibit the jarring, unusual phenomenology of strange cases such as the waterfall illusion. Fortunately, there is no reason to think that the stimuli used in the Mitroff study induced this sort of phenomenology. The results suggest instead that subjects had a clear, unambiguous percept of streaming objects and consistently indicated so.

The authors designed the stimuli to induce ‘a strong bias to consciously perceive streaming, by using smooth, constant, and reasonably fast motion’ (Mitroff et al. 2005: 74). They write that although ‘various manipulations can influence the overall proportion of bouncing vs. streaming responses, each individual event is always clearly seen as one or the other’ (op.cit.: 75n1). Moreover, the experience of bouncing in ambiguous bouncing/streaming displays is known to require other cues, such as disambiguating featural information. For example, using two distinct shapes can indicate that the object with a particular shape has changed trajectory after contact, causing a bouncing experience (Feldman & Tremoulet 2006). In the absence of such intervening cues, “the preference for straight motion paths dominates” and subjects experience streaming (Feldman & Tremoulet 2006: 131). The hypothesis that Mitroff et al.’s subjects had an unreported conscious experience of bouncing is therefore unmotivated. Accepting this hypothesis would require an ad hoc rejection of both the authors’ claim that their stimuli were successfully designed to rule out the experience of bouncing as well as the independent evidence that their stimuli lacked the features required to induce such an experience.

Finally, the stimuli are available online, where sceptical readers can view them for themselves.¹ Having watched the demonstrations several times (though not in a laboratory setting), I do not experience anything incoherent; rather, the two objects simply appear to stream through each other.

There is no plausible independent reason to think that the clear experience of the objects as streaming through each other is accompanied by any degree of conscious visual awareness of the objects as bouncing off each other. The evidence thus suggests that the object files were unconscious.

§3. Object files are genuinely perceptual

PPC provides reason for thinking that a content is represented unconsciously, not for thinking that it is genuinely perceived. For example, an illusory experience of size might fail to cohere with an unconscious dorsal-stream representation (Aglioti et al. 1995). But the vehicle of the latter content may be merely subpersonal. One might argue that object files are similarly subpersonal.

However, object files are a paradigm case of personal-level perceptual representation. Here are three reasons to think so.

First, object files are the outputs of constancy mechanisms. Phillips appeals to Burge's (2010) theory of perception in deeming constancy necessary for genuine perception (which, for Phillips, casts doubt on blindsight as evidence of unconscious perception [Phillips 2016: 434n20]). Object files track objects despite changes in their features and locations, thus exhibiting constancy across massive changes in retinal input.

Second, object files are used directly in guiding where we intentionally saccade. While involuntary eye movements may not constitute personal-level agency, moving one's gaze intentionally to track a desired object does. If subjects intend to saccade to an object and the object moves during the saccade, then subjects initiate a corrective saccade to fixate on the object in its new location; this corrective saccade accesses stored object files (Schut et al. 2017). The direct use of object files in such intentional actions suggests that they are personal-level.

Third, object files are available to working memory. Hollingworth and Rasmussen (2010) previewed colours in objects; then the objects moved; one second after the preview, test colours appeared in either their original objects or different objects. In a change-detection task, participants indicated whether any colours had changed. This task 'combined the reviewing

¹ <http://perception.yale.edu/Brian/demos/OF-BounceStream.html>

paradigm of Kahneman et al. (1992) with a colour memory task that is known to depend on VWM [visual working memory]' (Hollingworth & Rasmussen 2010: 546). Change-detection performance was enhanced when test colours appeared in their original objects. This experiment strongly suggests that object files are available to working memory.

Following Burge, Phillips requires evidence that perceptual states are available to 'central, coordinating agency' (Phillips & Block 2017: 166) for there to be reason to count them as personal-level. Working memory is a clear example of a central mechanism that coordinates between perception, cognition and action. The fact that object files are available to working memory shows that they are available to central, coordinating agency.

One might allow that object files available to working memory are personal-level, but question whether object files were available to working memory in the Mitroff study.² Perhaps some object files are personal-level (and available to working memory) and others are not. Mitroff et al. did not explicitly test working memory, so it is impossible to know with certainty whether the object files were available to working memory in their experiments.

However, there is reason to suspect that they were. In order to perform the object-reviewing task and indicate their phenomenology, subjects had to attend to the objects. Moreover, attention is required for consistent streaming phenomenology (Watanabe & Shimojo 1998). Attention is often thought to be a gateway for working-memory access (e.g., object-based attention increases performance on working-memory tasks [Scholl 2000]). Though the Mitroff experiments are unusual in that they involve ambiguous stimuli, they are not unusual in ways relevant to working-memory availability; in particular, subjects are not under cognitive load and attend to the objects without distraction.

Furthermore, the OSPB implicates working-memory storage. The ability to store past features of objects requires some form of short-term memory. The number of object files that can be deployed at once is around four (Pylyshyn & Storm 1988), which is roughly the capacity of working memory (Cowan 2001). Other short-term visual stores—iconic memory (Sperling 1960) and fragile visual short-term memory (Sligte et al. 2008)—have far higher capacity limits. This suggests the OSPB is achieved by object files stored in working memory. Since the OSPB was found in the Mitroff study, there is good reason to think the object files were available to working memory.³

² One could raise a similar worry about availability to guide saccades. Since availability to guide saccades appears to coincide with working memory storage (Schut et al. 2017), however, I focus here on working memory.

³ An anonymous referee questions whether object files themselves or merely some of the information they contain is meant to be personal-level. It is unlikely that merely features are stored in working memory, since they must be bound to representations of specific objects to generate the OSPB. Thus the object files themselves—

Even if there were no such evidence, however, it would not follow that the object files were subpersonal. Use by central, coordinating agency is not meant by Phillips to be a necessary condition on being personal-level. Instead, he claims only that if a representation is not useable for central, coordinating agency then there is no positive reason to count it as personal-level (Phillips & Block 2017: 181). If a representation is a token of a type that is ordinarily useable for central, coordinating agency, however, then there is at least some defeasible reason to think it is personal-level. One may object that this principle over-generates—e.g., representations of red can be stored in working memory, but that provides no reason to think every representation of red is personal-level. But unlike *red*, an object file is not merely a type of content that can be instantiated in both personal-level and subpersonal vehicles; it is a type of vehicle that plays a characteristic functional role in perception. The fact that object files are available to working memory in ordinary circumstances suggests that an object file is a personal-level type of mental representation, even if some tokens happen to occur outside working memory.

Object files are genuine personal-level perceptual representations. They can also occur unconsciously. Therefore, there is unconscious perception.

§4. Going forward

I hope to have shown that the specific example discussed above is a convincing case of unconscious perception. If not, PPC nonetheless provides philosophers and scientists working on this topic with a new methodological approach. There is no need to worry about whether subjects consciously detect a stimulus without being willing to report it, or to determine whether the subject's experience is 'no different from the subjective experience of nothing at all' (Peters et al. forthcoming).

Instead, researchers can (a) run experimental paradigms that are known to tap into genuine perceptual representations, (b) use stimuli that are ambiguous between two interpretations F and G but elicit conscious reports of F and (c) use a behavioural measure that indicates representation of G. If there is no reason to think that participants experience jarring

representations of objects as well as their features and locations—must be stored in working memory, which suggests that object files themselves are personal-level.

contradictory phenomenology, then the experiment provides evidence of unconscious perception.⁴

Bibliography

- Aglioti, S., DeSouza, J. F. X. and M. A. Goodale. 1995. Size-contrast illusions deceive the eye but not the hand. *Current Biology* 5(6): 679–85.
- Anaya, A. and S. Clarke. 2017. Naïve realism and unconscious perception: A reply to Berger and Nanay. *Analysis* 77(2), 267–273.
- Berger, J. and B. Nanay. 2016. Relationalism and unconscious perception. *Analysis* 76(4): 426–33.
- Blake, R. and N. K. Logothetis. 2002. Visual competition. *Nature Reviews Neuroscience* 3(1): 13–21.
- Block, N. 2016. The Anna Karenina principle and skepticism about unconscious perception. *Philosophy and Phenomenological Research* 93(2): 452–59.
- Burge, T. 2010. *Origins of Objectivity*. Oxford: OUP.
- Crane, T. 1988. The waterfall illusion. *Analysis* 48: 142–47.
- Feldman, J. and P. D. Tremoulet. 2006. Individuation of visual objects over time. *Cognition* 99: 131–65.
- Green, E. J. and J. Quilty-Dunn. Forthcoming. What is an object file? *British Journal for the Philosophy of Science*.
- Hollingworth, A. and I. P. Rasmussen. 2010. *Journal of Experimental Psychology: Human Perception and Performance* 36(3): 543–64.
- Kahneman, D., Treisman, A. and B. J. Gibbs. 1992. The reviewing of object files: Object-specific integration of information. *Cognitive Psychology* 24: 175–219.
- Luck, S., Vogel, E. and K. Shapiro. 1996. Word meanings can be accessed but not reported during the intentional blink. *Nature Scientific Reports* 383: 616–18.
- Marshall, J. C. and P. W. Halligan. 1988. Blindsight and insight in visuo-spatial neglect. *Nature* 336: 766–67.
- Mitroff, S. R., Scholl, B. J. and K. Wynn. 2005. The relationship between object files and conscious perception. *Cognition* 96: 67–92.

⁴ I am grateful to two anonymous referees for *Analysis* and to Ross Colebrook, E.J. Green, Zoe Jenkin, Eric Mandelbaum, Ian Phillips, Nic Porot, Nick Shea, Susanna Siegel and Tyler Wilson for comments on an earlier draft. Thanks also to audience members at the Institute of Philosophy for helpful discussion.

- Peters, M. A. K., Kentridge, R. W., Phillips, I. and N. Block. Forthcoming. Does unconscious perception really exist? Continuing the ASSC20 debate. *Neuroscience of Consciousness*.
- Peters, M. A. K. and H. Lau. 2015. Human observers have optimal introspective access to perceptual processes even for visually masked stimuli. *eLife*, doi:10.7554/eLife.09651.
- Phillips, I. 2016. Consciousness and criterion: On Block's case for unconscious seeing. *Philosophy and Phenomenological Research* 93(2): 419–51.
- Phillips, I. Forthcoming. Unconscious perception reconsidered. *Analytic Philosophy*.
- Phillips, I. and N. Block. 2017. Debate on unconscious perception. In *Current Controversies in Philosophy of Perception*, ed. B. Nanay, 165–92. New York: Routledge.
- Pylyshyn, Z. W. and R. Storm. 1988. Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision* 3(3): 179–97.
- Prinz, J. 2015. Unconscious perception. In *The Oxford Handbook of Philosophy of Perception*, ed. M. Mohan, 371–92. Oxford: OUP.
- Quilty-Dunn, J. (manuscript). Perceptual pluralism.
- Scholl, B. J. (2000). Attenuated change blindness for exogenously attended items in a flicker paradigm. *Visual Cognition*, 7(1/2/3): 377-96.
- Schut, M. J., Fabius, J. H., Van der Stoep, N. and S. Van der Stigchel. 2017. Object files across eye movements: Previous fixations affect the latencies of corrective saccades. *Attention, Perception, & Psychophysics* 79: 138–53.
- Sligte, I. G., Scholte, H. S. and V. A. Lamme. 2008. Are there multiple visual short-term memory stores? *PLoS One* 3(2), e1699. doi: 10.1371/journal.pone.0001699.
- Watanabe, K. and S. Shimojo. 1998. Attentional modulation in perception of visual motion events. *Perception* 27(9), 1041–54.