Acta Analytica (2025) 40:413–434
https://doi.org/10.1007/s12136-024-00611-5

Unconscious Perception, Action, and the Problem of Attribution

Paweł Jakub Zięba[1]

Received: 13 June 2024 / Accepted: 16 September 2024 / Published online: 2 October 2024

Abstract
According to Phillips, (1) genuine perception is attributable to the individual (i.e. it is a personal state/event, as opposed to sub-personal states/events in the individual's brain); (2) since unconscious perceptual representations are ill-suited to guide action, there is no good reason to attribute them to the individual; (3) not being attributable to the individual, they do not instantiate genuine perception, thereby failing to support the hypothesis that genuine perception can occur unconsciously. I argue that this reasoning is flawed and cannot be easily fixed. Phillips contends that unconscious perceptual representations do not guide action because they fail to meet certain conditions that are sufficient for action guidance. But those conditions may not be necessary for action guidance. Consequently, unconscious perceptual representations may guide action even if they do not meet Phillips' conditions. Furthermore, due to his commitment to the distinction between personal and sub-personal states/events, Phillips is not in a position to argue that his conditions are necessary for action guidance. For the distinction applies to action as well as to perception, and when genuine action is identified by personal-level criteria, Phillips' conditions turn out unnecessary for action guidance.

Keywords Unconscious perception · Action · Personal and sub-personal levels of explanation · Manifest and scientific kinds

1 Introduction

According to unconscious perception hypothesis (UP), 'episodes of the same fundamental kind as episodes of conscious perception can occur unconsciously' (Block & Phillips, 2017, p. 165). One of the main themes in the debate about UP is whether the putative instances of unconscious perception can guide action. For if they cannot, there is apparently no good reason to regard them as personal states rather than

✉ Paweł Jakub Zięba
pj.zieba@uj.edu.pl

[1] Jagiellonian University, Kraków, Poland

sub-personal states (Block & Phillips, 2017, p. 181; Phillips, 2018, pp. 493–494). Insofar as perception is supposed to be a personal state (i.e. a state attributable to the individual, as opposed to a state in the individual's brain), eligibility to guide action is a requisite for perceptuality.

Phillips (Phillips, 2018), the champion of scepticism about UP, calls this 'the problem of attribution'. The problem affects especially those of the putative instances of unconscious perception in which the stimulus is shown to influence behaviour and/or brain activity even though the subject's ability to discriminate that stimulus is at chance.[1] According to Phillips, in cases of this sort, unconscious perceptual representation of the stimulus is not available to the so-called Central Coordinating Agency (a placeholder for 'whichever subsystems subserve an agent's genuine, individual-level action' (Phillips, 2018, p. 497); from now on, CCA). Any behavioural reaction elicited by such unconscious representation is not a genuine action; it is merely an automatic reflex. Not being suitable to guide action, such representation is not attributable to the individual and thereby not a genuine perception.

In a reply to Phillips, Shepherd and Mylopoulos (Shepherd & Mylopoulos, 2021) have argued that unconscious perceptual representation *is* available to CCA (and thereby suitable to guide action) *if* 'availability to CCA' is specified as 'making a coherent contribution to the exercise of control'. For insofar as 'availability to CCA' is understood in this way, there is evidence that unconscious perceptual representation can guide action. It was found that behavioural effects associated with realising that one has committed an error can occur even if one is not consciously aware that one has made a mistake (Charles et al., 2013; Endrass et al., 2007; Fiscarella et al., 2019; Logan & Crump, 2010). Arguably, those behavioural effects result from unconsciously perceiving the features whose instantiation clashes with one's expectations.

Nevertheless, as Shepherd and Mylopoulos themselves acknowledge, 'availability to CCA' could be specified in some other way, leading to a different interpretation of the evidence. As we shall see, Phillips maintains that unconscious perceptual representations are not available to CCA because behavioural effects of unconscious perceptual representations fail to satisfy certain conditions that are sufficient for action guidance. Consequently, Phillips could respond that the evidence in question does not instantiate action guided by unconscious perceptual representation.

I will argue that this response does not work. My argument consists of two main points, which can be summarised as follows:

[1] Evidence for UP also includes cases where the subject's ability to discriminate the stimulus is above chance (even though the subject's report suggests lack of consciousness of the stimulus), but Phillips (Phillips, 2016) argues that this is most likely caused by unreported residual and/or transient conscious awareness of the stimulus. On this interpretation, the lack of report is due to a conservative criterion adopted by the subjects, who are inclined not to report perceiving the stimulus unless they have a certain amount of confidence that they did perceive it (hence 'the problem of criterion'). Cases of this sort may well involve a genuine perception guiding action, but it is doubtful whether that perception is unconscious. When the problem of attribution is considered jointly with the problem of criterion, it seems that 'the conditions for individually attributable perception suffice for perceptual consciousness' (Phillips, 2018, p. 501).

(1) Phillips claims that his conditions for action guidance are sufficient, but he does not insist that they are necessary. The idea is simple: when the conditions are met, there is a good reason to think that action guidance has occurred; when the conditions are not met, there is no good reason to think that. But, and this is my first point, since Phillips conditions are not meant to be necessary, it is completely unmotivated to say that when the conditions are not met, there is no good reason to think that action guidance has occurred. For there may be some other reason to think that action guidance has occurred. Given the circumstances in which a behaviour takes place, it may be perfectly reasonable to suppose that that behaviour is guided by some unconscious mental process.

(2) Let us suppose that Phillips actually insists that his conditions are necessary for action guidance. At first glance, this seems to rule out the possibility that action guidance has occurred even though Phillips' conditions are not met. But, and this is my second point, insisting that Phillips' conditions are necessary for action guidance renders unsound the reasoning on which the problem of attribution is based. To see why, notice first that that reasoning presupposes the distinction between personal/sub-personal states/events. That distinction applies to action as well as it does to perception, and when genuine action is identified by personal-level criteria, Phillips' conditions turn out unnecessary for action guidance. This means that Phillips cannot simultaneously hold that (X) his conditions are necessary for action guidance and that (Y) there is a distinction between personal/sub-personal states/events, because X is incompatible with Y. The necessity claim (which Phillips would have to make to resist my first point) is incompatible with the personal/sub-personal state/event distinction (which is a central assumption of the reasoning by which Phillips motivates the problem of attribution). Therefore, resisting my first point by insisting that Phillips' conditions are necessary for action guidance undermines the reasoning on which the problem of attribution is based.

The foregoing demonstrates that the reasoning on which the problem of attribution is based is flawed and cannot be easily fixed. For not only has Phillips failed to motivate the claim that the relevant evidence does not instantiate action being guided by unconscious perceptual representation, but also one of his central assumptions undermines his own reasons to doubt whether unconscious perceptual representation can guide action.

Section 2 takes a closer look at the reasoning by which Phillips brings up the problem of attribution. Section 3 shows what is wrong with that reasoning. Section 4 responds to two main objections against the central premise of my counterargument. Section 5 concludes.

2 A Closer Look at the Reasoning Behind the Problem of Attribution

This section scrutinises the reasoning behind the problem of attribution. Section 2.1 offers a reconstruction of that reasoning. Section 2.2 responds to some concerns about its exegetical accuracy.

2.1 The Reconstruction

Let us take a closer look at the way in which Phillips motivates the problem of attribution. The argument can be reconstructed as follows:

A1. Genuine action is subserved by Central Coordinating Agency (CCA).[2] [assumption]
A2. The presence of consciously intending, consciously deciding, consciously planning, or consciously willing to behave in a given way is a good reason to think that the behaviour in question is a genuine action (i.e. an output of CCA). [assumption]
A3. The behavioural effects of some[3] of the putative instances of unconscious perception are not accompanied by consciously intending, consciously deciding, consciously planning, or consciously willing to behave in the way one does. [assumption]
A4. There is a genuine and sharp distinction between perception qua personal state/event and perception qua sub-personal state/event. [assumption]
A5. Unconscious perceptual representation is unavailable to CCA. [explains A3 in the light of A1 and A2]
A6. Being unavailable to CCA, unconscious perceptual representation is not eligible to guide genuine action. [from A1, A5]
A7. There is no good reason to regard unconscious perceptual representation as a personal state rather than a sub-personal state. [from A4, A6]
A8. Unconscious perceptual representation does not instantiate genuine perception because genuine perception is a personal state. [from A4, A7]

A1 is just a paraphrase of Phillips' elucidation of CCA quoted above (Phillips, 2018, p. 497). Phillips does not say much about *how* CCA carries out genuine action. He only provides examples of cases where perceptual representation is unavailable to CCA despite influencing the subject's behaviour. Just because a perceptual

[2] Strictly speaking, Phillips (Block & Phillips, 2017, p. 169; Phillips, 2018, p. 494) only says that unconscious perceptual representation cannot guide action because it is unavailable to CCA. This leaves open the possibility that the availability to CCA is not necessary for guiding action because some other (i.e. non-perceptual) mental states can guide action despite not being available to CCA. But if that is the case, why cannot unconscious perceptual representation guide action in the way those other mental states do? Insofar as there is no good answer to this question (I do not think there is), the only way to make the argument convincing is to insist that the availability to CCA is necessary for guiding action.

[3] Those in which the stimulus discrimination is at chance (which makes it particularly plausible to say that perceptual representation of the stimulus is unconscious).

representation changes the reaction time, causes one's pupils to dilate, induces a skin conductance response, or even elicits attentional processing, it does not follow that it is available to CCA (Phillips, 2018, pp. 494–495). The lack of an explanation of the unavailability to CCA in terms of how CCA actually does its job is not necessarily a problem. For the purposes of the present discussion, it will suffice to conceive of CCA as a black box with an input and an output. Still, some explanation of what conditions have to be met in order for CCA to produce its output (i.e. genuine action) is necessary to make sense of Phillips' claim that unconscious perceptual representation is ineligible to be an input of CCA.

A2 and A3 are crucial components of Phillips' rationale for the claim that unconscious perceptual representation is unavailable to CCA. This is evidenced by two things Phillips says in published work. First, he claims that the availability of a mental state/event for guidance of intentional action constitutes a sufficient reason for the attributability of that mental state/event to an individual (Phillips, 2018, p. 497, 2020, p. 299).[4] Second, he maintains that a behavioural effect of unconscious perceptual representation does not instantiate action guidance because it is involuntary and/or inconsistent with the subject's knowledge, expectations, and intentions (Phillips, 2018, pp. 498–499).[5]

As the quoted passages demonstrate, Phillips expresses A3 directly: since the behavioural effects of unconscious perceptual representations are unintentional, involuntary, and unexpected, they are not accompanied by consciously intending, consciously willing, nor consciously planning to behave in a given way. A2, on the other hand, expresses an assumption underlying Phillips' assessment. By saying that (i) the behaviours in question are not outputs of CCA because they are not accompanied by consciously intending (willing…), Phillips effectively says that (ii) the presence of consciously intending (willing…) to behave in a given way is a sufficient reason to regard that behaviour as an output of CCA (see also footnote 4).

A4 is the thickest and presumably the most controversial premise in Phillips' reasoning. I accept it for the sake of the argument because I want to show that the reasoning behind the problem of attribution is flawed even if A4 is correct.

A4 applies to perception a certain general principle that specifies the distinction between personal and sub-personal levels of explanation (Dennett, 1969; Drayson, 2012). The latter distinguishes two types of psychological explanation. Personal explanations employ only folk-psychological concepts (e.g. belief, desire,

[4] '[T]he mere fact that a representation affects (e.g. speeds) the way that a subject responds is not a sufficient ground for attributing that representation to the individual as opposed to their brain or some psychological sub-system. The key (at least evidential) criterion for such individual attribution is, I propose, availability of the representation for guidance of intentional action' (Phillips, 2020, pp. 298–299).

[5] '[T]he effects of the subliminal words[…] are "radically uncontrollable" […], operating quite independently of the subject's own intentions. […] although the effects of the words may accidently coincide with the subject's intentions, their effects are non-volitional. Subjects cannot exploit the words to guide their behavior. The words merely affect their behavior outside their direct control' (Phillips, 2018, p. 498). '[T]he invisible prime directly activated its associated response regardless of the subjects' expectations and intentions. […] such activation bypasses the agent's own control and guidance as revealed by the fact that such activation occurs quite independently of the subject's knowledge and intentions' (Phillips, 2018, p. 499).

experience), whereas sub-personal explanations bring into play scientific concepts that pick various happenings in the nervous system (e.g. processing of sensory information, triggering muscle contractions due to increasing levels of neurotransmitters).

The point of the principle underlying A4 is that personal and sub-personal explanations do not share the same explanandum because they are about different types of states, events, or processes (personal and sub-personal, respectively). In other words, personal and sub-personal explanations have different truthmakers. Personal states/events are attributable to the individual (and not attributable to cognitive systems such as the perceptual system and its sub-systems), whereas sub-personal states/events are attributable to cognitive systems (and not attributable to the individual).

A4 applies the principle to perception, thereby drawing a sharp distinction between perception qua personal state/event (i.e. conscious perceptual experience) and perception qua sub-personal state/event (i.e. perceptual representation realised in the brain). The current debate about unconscious perception is structured in terms of this distinction. In particular, one of the standard assumptions of the debate is that genuine perception is an 'objective sensory representation attributable to the individual', where 'attributable to the individual' means 'a personal state' (Phillips, 2018, pp. 480–481, 490–491, 494).[6] This definition of perception, which comes from Burge (Burge, 2010, Chapter 9), is a common reference point not only in Phillips' work on unconscious perception, but also in the works of Phillips' opponents (see e.g. Block & Phillips, 2017, pp. 186–187; Peters et al., 2017, pp. 4–5; Shepherd & Mylopoulos, 2021, p. 3871). This is why the problem of attribution was introduced as a worry that unconscious perceptual representations are not individually attributable, even though the crux of the matter is their ability to guide action.

The distinction between personal and sub-personal states/events is 'genuine' in that it goes beyond the difference in the ways in which mental occurrences can be described. In other words, when we switch between sub-personal and personal levels of explanation (e.g. when we move from talking about information processing in the visual cortex to talking about what it is like to undergo a visual experience), we do not just switch between different ways of describing the same subject matter; we change the subject matter itself.

The distinction is 'sharp' in the sense that it is disjunctive and exhaustive. A mental occurrence can be either personal or sub-personal, but it cannot be both, and it cannot be neither. Also, a mental occurrence's being personal or sub-personal is not a scalar property that comes in degrees, nor is being personal or sub-personal relative to context.

A sub-personal state/event can be an enabling condition for a personal/state-event, but the former cannot be an essential feature of the latter. Since one cannot establish that X is an enabling condition for Y unless one already knows what Y is, that X is an enabling condition for Y presupposes that Y can be identified independently of X (i.e. that X is not among Y's essential features). Therefore,

[6] '[A] representation not attributable to the individual could not possibly be identified with perception in the ordinary sense which is plainly a personal or individual-level achievement. […] attempts to decouple perception and consciousness stumble at precisely this juncture' (Phillips, 2018, p. 481).

distinguishing perception qua personal state from perception qua sub-personal state entails that 'perceptual experience' is not co-extensive with 'perceptual representation' because the latter is just an enabling condition for the former.

Importantly, A4 *does not* boil down to the claim that there is a difference between perception qua personal state and perception qua sub-personal state. A4 also involves a story about *how* these two kinds of state differ from each other. In particular, A4 encompasses the claims that (a) genuine perception is a personal state and that (b) if a perceptual representation is unconscious and does not guide action, there is no good reason to consider it a personal state (i.e. to attribute it to the individual).

In the current debate about unconscious perception, it is commonly assumed that A4 entails (a) (Block & Phillips, 2017, p. 181; Burge, 2010, p. 374; Peters et al., 2017, pp. 7–8; Phillips, 2018, pp. 480, 482). In particular, the idea that genuine perception is not (or may not be) attributable to an individual (i.e. a personal state) is incompatible with Burge's definition of genuine perception, which is assumed by Phillips and his opponents in the debate about unconscious perception.

The idea that perception qua personal state is not (or may not be) able to guide action (i.e. the denial of (b)), while intelligible, is rather counterintuitive. Strictly speaking, Phillips only says that the ability of a perceptual representation to guide action is *sufficient* for considering that representation a personal state; he does not claim that it is *necessary*. But he also claims that if unconscious perceptual representation is unable to guide action, we have no reason to attribute it to the individual, i.e. categorise it as a personal state (Block & Phillips, 2017, p. 181). So, whether the ability to guide action is considered necessary or sufficient for being a personal state, Phillips' view is that unconscious perceptual representation is not a personal state unless it is able to guide action (i.e. unless it is available to CCA).

One may question A4 by disputing the assumption that the distinction between personal and sub-personal levels of explanation maps onto a corresponding distinction between personal and sub-personal states/events (Drayson, 2012). However, Phillips (Phillips, 2018) can drop the personal vs. sub-personal distinction and instead formulate A4 in terms of an independent (albeit somewhat related) distinction between manifest and scientific kinds. According to the latter, perception in its ordinary sense is a manifest kind, i.e. 'a kind whose instances we identify and re-identify on the basis of their manifest properties' (Johnston, 1997, p. 565; Phillips, 2018, p. 477). Perceptual experience instantiates perception qua manifest kind because we identify and re-identify it in relation to its phenomenal character. Unconscious perceptual representation, in contrast, lacks phenomenal character. It is a scientific kind, in that it can only be identified and re-identified by third-personal methods. Perception qua scientific kind might be an enabling condition for perception qua manifest kind, but no perception qua manifest kind is enabled until perceptual experience occurs. Now, even if the personal vs. sub-personal states/events distinction is flawed, Phillips can argue that perceptual representation must be available to CCA in order to enable perception qua manifest kind.

While the foregoing specification of A4 could be questioned, it faithfully represents Phillips' position, and I am assuming it for the sake of the argument because I want to show that Phillips' argument fails even if this specification is correct.

A5 explains A3 in the light of A1 and A2. The explanation in question is simple. We are assuming that all behaviours that count as genuine actions are outputs of CCA (A1) and that a behaviour can be considered an output of CCA if it is accompanied by consciously intending (willing…) to behave in the way one does (A2). The phenomenon to be explained is the fact that behavioural effects of unconscious perceptual representations fail to satisfy this condition (A3). Therefore, our explanandum is the fact that a good indicator of a behaviour's being an output of CCA is absent from behavioural effects of unconscious perceptual representations. Phillips explains this absence with the hypothesis that unconscious perceptual representations are not available to CCA (A5).[7] Obviously, A5 does not follow from A1 and A2 because one could offer a different explanation of A3 (e.g. that behavioural effects of unconscious perceptual representations are accompanied by unconscious planning). But this does not change the fact that A5 is one of the simplest possible explanations of A3, given A1 and A2.[8]

If genuine action is subserved by CCA (A1) and A5 correctly explains A3 (i.e. unconscious perceptual representations are not available to CCA), then unconscious perceptual representations are unable to guide action (A6). The behavioural effects of unconscious perceptual representations are not outputs of CCA. They do not instantiate action guidance.[9]

Once it is assumed that A4 encompasses (a-b), it becomes clear that the conjunction of A4 and A6 entails A7. If unconscious perceptual representation has to be able to guide action in order to instantiate perception qua personal state (given A4),

[7] 'Where attentional responses are completely stimulus-driven reflexes, operating entirely outside of voluntary control […], and possibly mediated by subcortical pathways […], I am unpersuaded that we must think of them as exercises of individual-level agency' (Block & Phillips, 2017, p. 181). Here the 'attentional responses' are behavioural effects of unconscious perceptual representation. Phillips (Phillips, 2018, p. 497) also cites Frankfurt: 'When we act, our movements are purposive […] their course is guided […] The dilation of the pupils […] does not mark the performance of an action by the person; his pupils dilate, but he does not dilate them. This is because the course of the movement is not under his guidance. The guidance in this case is attributable only to the operation of some mechanism with which he cannot be identified' (Frankfurt, 1988, p. 73). According to Phillips, behavioural effects of unconscious perceptual representation are like pupil dilation because the former are non-intentional, involuntary, unexpected. This comparison illustrates the intuition that unconscious perceptual representations are unavailable to CCA. Therefore, A5 explains A3 in the light of A1 and A2.

[8] An anonymous referee has suggested that Phillips overdramatises the differences between conscious and unconscious perceptual episodes, to the extent that it becomes hard to see that these two types of episodes could have anything in common. Perhaps the source of this problem resides in Phillips' positive view about perception. According to the latter, perception in the ordinary sense is a primitive (i.e. psychologically unanalysable) relation of *conscious* acquaintance to aspects of mind-independent reality. The occurrence of this relation is caused and underlied by computational processes involving subpersonal representations in the brain, but the relation itself is not representational (French & Phillips, 2023). This conception seems to beg the question against UP by suggesting that perception is inherently conscious. That said, one may also wonder whether a certain amount of perceptual processing, significant but insufficient for the occurrence of conscious acquaintance, could cause and underlie a relation of *unconscious* acquaintance, i.e. an unconscious instance of perception in the ordinary sense (Anaya & Clarke, 2017; Zięba, 2019).

[9] 'Subjects cannot exploit the words to guide their behavior. The words merely affect their behavior outside their direct control' (Phillips, 2018, p. 498). '[I]nvisible primes are not useable or exploitable by the individual to guide their actions' (Phillips, 2018, p. 499).

and unconscious perceptual representation is ineligible to guide action (A6), then unconscious perceptual representation is not perception qua personal state (A7).[10]

The acknowledgment that A4 encompasses (a-b) also secures the move from A4 and A7 to A8. If genuine perception is a personal state (given A4), and unconscious perceptual representation is not a personal state (A7), then unconscious perceptual representation is not genuine perception (A8).[11]

2.2 Is the Reconstruction Accurate?

One might question the exegetical accuracy of my reconstruction by pointing out that Phillips can allow that intending (willing…) to behave in a given way can sometimes be unconscious. But this is beside the point, because Phillips cannot allow that unconsciously intending (willing…) to behave in a given way can play the same role in guiding action as consciously intending (willing…) to behave in that way. What blocks the move in question is A4, which is an indispensable component of the argument.

If one applies the principle behind A4 to perception, consistency requires one to take the same approach to other mental occurrences, including intending (willing…). If there is a difference in kind between perception qua personal state/event and perception qua sub-personal state/event, the same applies to intending (willing…). So even if Phillips allows for, say, unconsciously deciding, consistency requires him to regard it as sub-personal. On this view, unconsciously deciding is a sub-personal state/event, a mere enabling condition for deciding qua personal state. As such, unconsciously deciding cannot play the same role in guiding action as consciously deciding. For if unconsciously deciding can play the same role in guiding action as consciously deciding, why think that unconsciously perceiving cannot play the same role in guiding action as consciously perceiving?

To justify any such asymmetry, one would have to (i) identify a feature that enables deciding to play its specific role in action guidance independently of consciousness and (ii) show that perceiving does not possess that kind of feature. Let us call that feature 'agentive consciousness-independence' (ACI for short). To say that X is ACI means that the specific role that X plays in action guidance does not depend on consciousness. So if a mental episode M (e.g. deciding, perceiving) is ACI, then M can realise its specific role in guiding action whether or not M is conscious. Thus understood, ACI is a structural feature that could be (at least in principle) coherently attributed to both perceiving and deciding, even though the roles that perceiving and deciding play in action guidance are different.

[10] '[W]hen a representation is unavailable to central agency, we lack a positive ground for attribution. […] If they [attentional responses caused by unconscious perceptual representations] are not [exercises of CCA], we lack positive reason for thinking of the perceptual representations […] as constituting individual-level perception' (Block & Phillips, 2017, p. 181).

[11] '[A] representation not attributable to the individual could not possibly be identified with perception in the ordinary sense which is plainly a personal or individual-level achievement' (Phillips, 2018, p. 481). 'If perception is by the individual and such attribution requires availability for action guidance, then this is not perception' (Phillips, 2018, p. 498).

Now, my point is that there seems to be no good reason to believe that deciding is ACI but perceiving is not ACI. In particular, it is hard to imagine a reason for which deciding would be ACI but perceiving would not. I am not saying that this is logically impossible, but I see no good reason why this could be the case. Just because unconsciously deciding would presumably be a part of CCA's output (whereas unconscious perceptual representation is not an output of CCA) does not make unconsciously deciding eligible to play the same role as deciding qua personal state. Personal states are identified from the first-person perspective, and there is nothing it is like to unconsciously decide. Besides, even if unconsciously deciding was granted the status of 'genuinely deciding' after all (which is, in all likelihood, simply incompatible with A4), Phillips would be hard-pressed to explain why on earth are unconscious perceptual representations incapable of causing unconscious decisions (which he really would not want to allow, since that would render unconscious perceptual representations capable of guiding genuine action).

It seems that whatever reason is given for the claim that unconsciously intending (willing…) to behave in a certain way is a personal state/event, or can play the same role as a personal state/event, it is going to cast serious doubt on Phillips' case for the asymmetry between conscious perception and unconscious perceptual representation. If unconsciously intending (willing…) to behave in some way is considered a personal state/event, A4 starts to work against the sceptic about UP, because the standards for being a personal state/event become so low that it is now unclear why unconscious perceptual representation fails to meet them. For there is no evident difference between (i) unconscious perceptual representation and (ii) unconscious intending (willing…) that would justify placing (i) and (ii) on different sides of the personal/sub-personal divide. If (ii) counts as a personal state/event, why (i) does not?

Admittedly, the personal/sub-personal divide does not always overlap with the conscious/unconscious divide. One can consistently hold that unconscious belief is a personal state whereas unconscious perceptual representation is not. However, categorising unconscious decision as personal undermines Phillips' point that the content of unconscious perceptual representation is unavailable for personal-level decision-making. If deciding qua personal state can only occur consciously, it is reasonable to say that it can only be guided by conscious contents. But if deciding qua personal state can be unconscious, it is unclear why it could not be guided by unconscious contents (e.g. contents of unconscious perceptual representations).

If, on the other hand, unconsciously intending (willing…) to behave in a certain way is considered sub-personal but functionally equivalent to consciously intending (willing…) to behave in that way, A4 is violated, or at least devoid of any significance. For if personal states/events can be functionally equivalent to sub-personal states/events, what is the point of drawing the distinction in the first place? Such functional equivalence would undermine the personal/sub-personal distinction because it entails that attributability to the individual is not settled by functional role.

This is why the sceptic about UP has to insist that unconsciously intending (willing…) to behave in a certain way, if there are such things, are irrelevant to the reasoning behind the problem of attribution. Given A4, they are ineligible to be features

by which we identify and re-identify instances of genuine action. This is also the reason why A2 specifies those features as conscious.

Assaf Weksler pointed out to me that my reconstruction omits another reason to think that unconscious perceptual representation is unavailable to CCA. That reason is this: the behavioural effects of at least some cases of unconscious perceptual representation (most notably subliminal priming) are too weak to instantiate action guidance. Priming might activate a concept, which can result in a faster (or slower) stimulus identification in a subsequent phase of the trial, but this does not mean that the subliminal prime is available for decision-making.

For example, a mere unconscious perceptual representation of a backward-masked tomato will not yield a decision (neither conscious nor unconscious) to grab the tomato, even if the subject actually wants a tomato. In this sense, the tomato is not available for (conscious or unconscious) decision-making. All that the masked tomato does is activating the concept 'tomato' in the subject's brain. This can make the subject notice tomatoes faster later on, or it can bias their thinking in the direction of tomatoes.

It seems that this is what Phillips had in mind when he wrote that

> 'to the extent that subjects cannot themselves use representations of the attended but unseen objects to guide their responses, such representations do not witness genuine perception. The differential processing they produce is instead akin to a stimulus-driven reflex, operating entirely outside of voluntary, agentive control.' (Phillips, 2018, pp. 495–496)

Consider the subject S who is having an unconscious perceptual representation R. The idea at hand is that R is unavailable to S's CCA because R's being unconscious precludes S from being able to form any intention (volition…) regarding whatever it is that R represents. At best, R can influence the formation or execution of the intentions (volitions…) that S forms with respect to the contents of S's conscious perceptual representations.

The problem with this reasoning is that it presupposes that R can only guide an action if R initiates that action. If so, the behavioural effects of many putative instances of unconscious perception fail to instantiate action guidance. However, it seems perfectly coherent that a perceptual representation can guide action (or participate in guiding action) even if it is by itself insufficient to cause that action to occur. If so, making a behaviour faster, or slower, or more efficient, or less efficient, may well count as guiding action.

Consider the following simple condition for participation in guiding action. A perceptual representation R participates in guiding an action A if the content of R at least partially rationalises (i) the occurrence of A and/or (ii) the way A was executed. This simple condition is met by at least some cases of unconscious perceptual representation discussed in the literature. For example, attentional effects of an unconscious perceptual representation are rationally explicable in terms of the content of that representation (Jiang et al., 2006). As Watzl points out,

> 'the attention effects of masked cues just like the attention effects of non-masked cues appear to be contingent on the subject's overall intentions and

> goals. We see the same form of motivational penetration, or contingent capture, as in cases of conscious attention guidance. They are integrated with the (whole) subject's motivational system in a way that we would expect if the reaction to the mask was genuinely perceptual, but not if it were subsubject level.' (Watzl, 2017, p. 122)

Unless the UP-sceptic presents some good reason to believe that the simple condition is mistaken, unconscious perceptual representation can be available to CCA even if it never initiates action. As far as I can tell, the simple condition is perfectly viable. Moreover, given what I say below about 'not consciously intended actions' (Sect. 3), unreliability of introspection (Sect. 3), and 'intelligent reflexes' (Sect. 4), the condition of initiating action seems unduly restrictive.

My simple condition, just as the more sophisticated condition offered by Shepherd and Mylopoulos (Shepherd & Mylopoulos, 2021, p. 3873) which I mentioned in the Introduction, could be questioned. As Shepherd and Mylopoulos (Shepherd & Mylopoulos, 2021, p. 3889) observe, 'on some explications of attributability, or of availability to central coordinating agency, an unconscious state qualifies, while on other explications, the same state fails to qualify.' But as long as the problem of attribution hinges on an optional and controversial account of action guidance, its significance is severely limited, since the UP-enthusiast can simply choose an account of action guidance that suits them and move on.

If the condition of initiating action is rejected as too restrictive, A2 and A3 are the only motivation left for A5 (to reiterate, A5 is Phillips' explanation of why A3 is the case, given A1 and A2). This is why A2 and A3 are crucial components of the reasoning behind the problem of attribution.

3 What Is Wrong with the Reasoning Behind the Problem of Attribution?

As mentioned above, A2 is absolutely crucial for the reasoning behind the problem of attribution. But the most important thing about A2 is not what it says, but what it *does not* say. Namely, it *does not* say that the presence of consciously intending (willing…) to behave in a given way is *necessary* for genuine action. Although in published work Phillips is not entirely clear about this, in conversation he told me that the argument does not rely on any such necessity claim. The idea is rather that insofar as behavioural effects of unconscious perceptual representations are not accompanied by consciously intending (willing…) to behave in a given way, there seems to be no reason to regard those behaviours as genuine actions. Since the presence of consciously intending (willing…) to behave in a given way is sufficient for action guidance (and thereby marks a genuine action), the UP-enthusiast can validate UP by adducing evidence that a behaviour elicited by unconscious perceptual representation is accompanied by consciously intending (willing…) to behave in that way. But unless such evidence is provided, there is no good reason to believe that unconscious perceptual representation can guide action and thereby also no reason for attributing such representations to the

individual. Unless it is shown that unconscious perceptual representation initiates a consciously intended behaviour, or at least a consciously intended modification of a behaviour, the impact of unconscious perceptual representation on behaviour is better viewed as an involuntary reflex that falls short of action guidance.

This argument is unpersuasive. Since Phillips does not insist that consciously intending (willing…) to behave in a given way is necessary for genuine action, the UP-enthusiast is free to dismiss the problem of attribution as a pseudo-problem. For insofar as genuine action can occur without consciously intending (willing…) to behave in a given way, the behavioural effects of unconscious perceptual representations can instantiate action guidance even if they are *never* accompanied by consciously intending (willing…) to behave in the way one does. For all that Phillips has told us, unconscious perceptual representation may guide an action even if its influence on that action is not reflected in (or even incompatible with) one's conscious decisions or intentions and even if it does not initiate that action. Just because something is *sufficient* for action guidance does not automatically make it *necessary* for action guidance. Therefore, the fact that behavioural effects of unconscious perceptual representation fail to satisfy the conditions identified by Phillips as *sufficient* for action guidance is *not* a reason to believe that unconscious perceptual representation does not guide action (i.e. that it is unavailable to CCA).

Given the foregoing, the problem with Phillips' argument is not just that A5 is not entailed by A3, but also that A3 is not a good reason to believe that A5 is true. Because the move from A3 to A5 lacks compelling motivation, the UP-enthusiast is free to reject A5. For example, they may respond that the fact that unconscious perceptual representations can causally influence behaviour in rationally explicable ways (see e.g. Jiang et al., 2006) suffices to regard such representations as eligible to guide action. Whether this reply is correct or not, *nothing* in Phillips' argument undermines it.

What if Phillips actually did insist that intending (willing…) to behave in a given way is necessary for that behaviour to count as genuine action? Would that compel the UP-enthusiast to search for evidence that behavioural effects of unconscious perceptual representation can meet this requirement?

Of course, the UP-enthusiast could simply reject the requirement. For example, they could respond that unconscious perceptual representation guides action by causing the subject to *unconsciously* intend (will…) to behave in a given way. That reply would most likely turn the debate about unconscious perception into a debate about necessary and sufficient conditions for genuine action. But the fact of the matter is that Phillips simply *cannot* amend his argument by adding the necessity requirement to A2. The goal of the remainder of this paper is to explain why this is so. Let us start by considering the amended version of the argument:

A1. Genuine action is subserved by Central Coordinating Agency (CCA). [assumption]
A2*. Consciously intending, consciously deciding, consciously planning, or consciously willing to behave in the way one does is *necessary* for that behaviour to count as genuine action (i.e. as an output of CCA). [assumption]

A3. The behavioural effects of some[12] of the putative instances of unconscious perception are not accompanied by consciously intending, consciously deciding, consciously planning, or consciously willing to behave in the way one does. [assumption]
A4. There is a genuine and sharp distinction between perception qua personal state/event and perception qua sub-personal state/event. [assumption]
A5. Unconscious perceptual representation is unavailable to CCA. [explains A3 in the light of A1 and A2*]
A6. Being unavailable to CCA, unconscious perceptual representation is not eligible to guide genuine action. [from A1, A5]
A7. There is no good reason to regard unconscious perceptual representation as a personal state rather than a sub-personal state. [from A4, A6]
A8. Unconscious perceptual representation does not instantiate genuine perception because genuine perception is a personal state. [from A4, A7]

At first glance, it seems that replacing A2 with A2* fixes the argument. If genuine action (i.e. the output of CCA) cannot occur without consciously intending (willing…) to behave in a given way (as A2* has it), then A3 constitutes a good reason to believe that A5 is true. Although A3 still does not entail A5, A2* and A3 seem to make it quite clear that behavioural effects of unconscious perceptual representations are not genuine actions.

Nevertheless, replacing A2 with A2* moves the problem of attribution from the frying pan into the fire. For, as I argue below, A2* is incompatible with A4, which renders the reasoning behind the problem of attribution unsound. I start by presenting my argument in a schematic form and then elaborate on the key claims. Here is the argument:

A2*. Consciously intending, consciously deciding, consciously planning, or consciously willing to behave in the way one does is *necessary* for that behaviour to count as genuine action (i.e. as an output of CCA). [an assumption of the amended reasoning behind the problem of attribution]
A4. There is a genuine and sharp distinction between perception qua personal state/event and perception qua sub-personal state/event. [an assumption of the amended reasoning behind the problem of attribution]
B1. There is a genuine and sharp distinction between personal and sub-personal states/events. [presupposed by A4]
B2. There is a genuine and sharp distinction between action qua personal state/event and action qua sub-personal state/event. [a consequence of B1]
B3. There are cases of action qua personal state/event that do not involve consciously intending, consciously deciding, consciously planning, or consciously willing to behave in the way one does. [from B2]

[12] Those in which the stimulus discrimination is at chance.

B4. Consciously intending, consciously deciding, consciously planning, or consciously willing to behave in the way one does is *not* necessary for that behaviour to count as genuine action (i.e. as an output of CCA). [from B3]
B5. A4 entails that B4 is true [from B1, B2, B3], whereas A2* entails that B4 is false.
B6. A2* and A4 cannot both be true. [from B5]
B7. The problem of attribution rests on unsound reasoning. [from B6]

B1 follows from A4 because the latter is just an application of the general principle expressed by B1 to perception. To be clear, B1 would not follow from A4 if A4 was considered in isolation from the context described in Sect. 2.1. But that context is crucial for understanding where A4 came from. A4 is not just a view about perception; it is an application of a more general principle expressed by B1. If A4 was considered in isolation from B1, it would be completely ad hoc and unmotivated. Hence, the commitment to A4 presupposes the commitment to B1. The entailment from A4 to B1 holds irrespective of whether the reasoning behind the problem of attribution is specified in terms of the personal/sub-personal state/event distinction or in terms of the manifest/scientific kind distinction. On the latter reading, A4 claims that there is a genuine and sharp distinction between perception qua manifest kind and perception qua scientific kind, while B1 is a more general principle stating that there is a genuine and sharp distinction between manifest kinds and scientific kinds.

What exactly makes a state/event personal is a matter of an ongoing debate (see e.g. Drayson, 2012; Westfall, 2024). Still, whatever criterion is adopted, it is going to pertain to the personal level of explanation, which is shaped by (i) the way the world manifests to us in conscious experience and (ii) folk-psychological platitudes characterising this manifest image.[13]

Suppose that either the personal/sub-personal state/event distinction, or the manifest/scientific kind distinction (or both), accurately expresses how things are (i.e. that A4 is true). This means that our most basic understanding of what perception is comes from perception qua personal/manifest kind state/event. In other words, A4 entails that perception in the ordinary sense of 'perception' is individuated in relation to (i) its phenomenal character (which is characterised by what it is like to undergo a perceptual experience) and (ii) folk-psychological platitudes characterising that character (e.g. that perception relates one to items in one's environment).

Given B1, the same story applies to all mental occurrences.[14] If one applies the principle expressed by B1 to perception, consistency requires one to take the same approach to other mental occurrences, including action. I see no reason why action

[13] As illustrated by Dennett's classic discussion of pain (Dennett, 1969, p. 91).

[14] One might object that the personal/sub-personal state/event distinction and the manifest/scientific kind distinction apply to some mental occurrences but not to others or that they apply exclusively to perception. Since these are genuine logical possibilities, B2 is not a logical consequence of B1. Nevertheless, it is hard to see what reason could be given for such a restriction other than the need to block the argument presented here. Unless the proponent of B1 provides such a reason, denying B2 creates an unaccounted incongruity in their position.

should be treated differently than perception in this case. If perception in its ordinary sense is a manifest kind, the same applies to action. Therefore, Phillips has to identify genuine action in the same way he identifies genuine perception. The commitment to A4 compels him (via B1) to identify genuine action in personal-level/manifest terms: by action's specific phenomenology, and by the way in which people think and talk about action in everyday situations (hence B2).

Given B1, personal states/events cannot be identified by sub-personal states/events, just as states/events of manifest kind cannot be identified by states/events of scientific kind (see Sect. 2.1). Any state/event of CCA is by definition sub-personal. The availability of a perceptual representation to CCA is a state/event of scientific, not manifest, kind. Hence, the availability to CCA is ineligible to be an essential feature of action qua personal/manifest kind state/event. By contrast, consciously deciding to behave in the way one does is eligible to be such a feature because it is a personal/manifest kind state/event. But unconsciously deciding to behave in the way one does is also ineligible to be an essential feature of action in the ordinary sense. At best, it might be a requisite for action qua sub-personal/scientific kind state/event.

One more thing about B2. Throughout the paper, by 'genuine action' I just mean something that one does, as opposed to something that happens to one. One might object that a narrower sense of 'action' is relevant here. As mentioned above, Phillips (Phillips, 2020) says that the availability of a mental state/event for guidance of intentional action constitutes a sufficient reason for the attributability of that mental state/event to an individual. However, it is a matter of debate whether all action is intentional, and my argument is neutral about this.

Now, if B2 is true, B3 is true as well. In particular, when genuine action is identified in personal-level/manifest terms, it turns out that neither consciously intending, nor consciously deciding, nor consciously planning, nor consciously willing to behave in the way one does (nor any other event of this sort) is necessary for genuine action.

To see why, compare tapping one's foot with being pushed and falling on the ground. The former is usually unintentional (or at least usually not consciously intended), but it is nonetheless something one does, not something that happens to one. The latter, by contrast, is something that happens to one, not something one does.

To be sure, tapping one's foot is notably less sophisticated than many mundane activities, like walking up the stairs. Still, tapping one's foot and walking up the stairs are in an important sense alike. Namely, they are alike in the sense in which leaning forward when the bus suddenly stops and walking up the stairs are *not* alike. Tapping one's foot, just like walking up the stairs, is not like leaning forward when the bus suddenly stops, because it is something one does, not something that happens to one. By the same token, being pushed and falling on the ground is like leaning forward when the bus suddenly stops, because it is something that happens to one, not something one does.

If tapping one's foot is something one does even though it is usually not consciously intended, then it constitutes a case of behaviour that (i) is not accompanied by consciously intending (willing…) to behave in the way one does, and yet (ii) passes the most basic test for genuine action. If tapping one's foot is a genuine

action, it is attributable to the individual. Assuming that all cases of genuine action involve some kind of action guidance, and that guiding action is a marker of personal states, tapping one's foot is (usually) guided by some unconscious personal state/event. Whatever that guiding state/event is, it is attributable to the individual.

Call cases such as tapping one's foot *not consciously intended actions* (from now on, NCIs). B3 claims that NCIs meet the standard for action qua personal/manifest kind state/event even though they do not involve consciously intending (willing…) to behave in the relevant way.

NCIs are common in everyday life. Apart from tapping one's foot, examples of NCI include humming, fidgeting, repositioning oneself in a chair, idly drumming one's fingers on the table, unawares moving one's tongue in one's mouth while driving a car. Arguably, the category of NCI also includes *arational actions* (Hursthouse, 1991), which occur when one acts out of emotion, without the involvement of any deliberation (e.g. kissing or lightly touching in passing a photograph of a person one loves, destroying or damaging anything remotely connected to a person one hates, jumping or running out of excitement, hiding one's face out of fear, talking or posturing to oneself in the mirror out of pride).

NCIs are a matter of controversy. Some philosophers consider NCIs intentional actions (O'Shaughnessy, 2008), whereas others mention NCIs in support of the view that no mental activity is essential for action (Runyan, 2014). Either way, NCIs are actions (i.e. things one does, not things that happen to one), and yet no consciously intending (willing…) to behave in the relevant way is necessary for them to occur. This is not to deny (nor to assert, for that matter) that conscious proprioception of the movement is a requisite for action (O'Shaughnessy, 2008, p. 356). The point is that even if NCIs are intentional, the subject does not consciously intend (will…) to behave in the way they do (cf. Rosenthal, 2008). Even if NCIs do involve intending (willing…) to behave in the relevant way, the subject does not *consciously* intend (will…) to behave in that way, which means that the requirement expressed by A2* is not met.

Denying the existence of NCIs has problematic consequences. Arguably, one usually does not consciously intend (will…) to tap one's foot, to move one's tongue in one's mouth while being focused on driving a car, to fidget on a chair. More often than not, one finds out that one is doing such things long after one started doing them. And yet the rejection of NCIs predicts that in every such case the behaviour is accompanied by consciously intending (willing…) to behave in that way. This prediction seems phenomenologically inadequate. It mischaracterises the way we experience our own behaviour in everyday situations.

Rejecting NCIs is also at odds with well-established evidence for the unreliability of introspection, which suggests that people often do not know the reasons behind their own decisions and actions (Nisbett & Wilson, 1977; Schwitzgebel, 2008; Wilson, 2002). If NCIs do not exist because every possible candidate for an NCI is accompanied by conscious intending (willing…) to behave in the relevant way, we should expect that people are rarely ignorant of the reasons behind their own decisions and actions. But evidence from psychology suggests that this is not the case.

That said, it also seems clear that people can be held responsible for their NCIs, at least in the sense of authorship of the behaviour. Just because one does not know

why one did something does not mean that one did not do it. People do not tap their feet because they were forced to do so, nor does it happen to them by accident (at least not usually). It is something *they* do, not something that happens to them.

B3 is the key premise of my argument, and I provide some additional support for it in Sect. 4, where I address two main objections against it. But let us finish the argument first. B4 draws the conclusion from B3: consciously intending (willing…) to behave in the way one does *is not* necessary for that behaviour to count as genuine action (i.e. as an output of CCA). If B4 is true, the rest of the argument follows smoothly.

The possibility of NCI entails that A2* is false. Insofar as consciously intending (willing…) to behave in the way one does is completely absent from NCI, the behavioural effect of unconscious perceptual representation can be a genuine action even if it is not accompanied by consciously intending (willing…) to behave in a given way.

On top of that, and to reiterate what I have already explained in Sect. 2.2, any attempt to save A2* by suggesting that NCIs involve *unconsciously* intending (willing…) to behave in the way one does would be futile. For to save A2*, Phillips would have to allow that unconsciously intending (willing…) to behave in the way one does can play the same role in guiding action as consciously intending (willing…) to behave in the way one does. In particular, he would have to reformulate A2* by saying that *either consciously or unconsciously* intending (willing…) to behave in the way one does is necessary for that behaviour to count as genuine action. The problem with this move is that it would undermine Phillips' own reason to think that a sub-personal/scientific kind state/event cannot play the same role in guiding action as a personal/manifest kind state/event. For example, Phillips would have to allow that unconsciously deciding (a sub-personal/scientific kind state/event) can play the same role in guiding action as consciously deciding (a personal/manifest kind state/event), which would violate his own reason to believe that unconscious perceptual representation (a sub-personal/scientific kind state/event) cannot play the same role in guiding action as conscious perception (a personal/manifest kind state/event). Simply put, there is no way to save A2* without violating A4.[15]

4 Two Objections Against B3

One might try to resist my argument by suggesting that NCIs are consciously intended because they involve residually and/or transiently conscious intending (willing…) that escapes consciousness too quickly to be reportable.[16] This seems possible, but not all putative examples of NCI can be explained in this way. The explanation in question cannot account for 'intelligent reflexes', i.e. cases where the

[15] Recall the discussion of ACI in Sect. 2.2. If intending (willing…) is ACI, why perceiving is not ACI? I see no good reason why.

[16] Notice that this reply effectively reduces the problem of attribution to just another form of the problem of criterion (see footnote 1).

right action in a given context is unconsciously selected and initiated 'with rapid retrieval and execution of the appropriate control policy or automatised response' (Krakauer, 2019, pp. 827, 829).

According to Krakauer (Krakauer, 2019, p. 829), intelligent reflexes are involved in performance of any motor skill, which makes them a common phenomenon. Furthermore, intelligent reflexes neatly explain the greatness of sports champions. Attributing an excellent performance to an athlete (and not to habits or instincts of that athlete) is not a mistake precisely because a behaviour can be voluntary and intelligent even if in some respects it is also automatic. Again, just because you do not consciously intend (will…) to behave in a given way does not mean that it is not you who is behaving in that way or that the behaviour is somehow done for you by your body.

Instead of insisting that NCIs involve consciously intending (willing…) to behave in a given way, the critic could argue that the absence of consciously intending (willing…) to behave in the way one does indicates that NCIs lack proprietary agentive phenomenology and for this reason fall short of genuine action. However, this objection is also unconvincing.

According to Mylopoulos and Shepherd (Mylopoulos & Shepherd, 2020), 'a phenomenal property is proprietarily agentive if it is instantiated by a state or process central to *some* actional episodes and not instantiated by states and processes that occur outside of actional episodes'. They enumerate six aspects characterising agentive phenomenology: *purposiveness* (a sense of mandating the behaviour), *mineness* (a sense of authorship of the behaviour), *execution* (a sense of accomplishing a goal), *action perception* (a perceptual experience of various events that are related to the behaviour), *action assessment* (any experience related to evaluation of the behaviour), *freedom* (a feeling that the behaviour is free).

NCIs clearly instantiate *mineness* and *action perception* (it is me who is tapping my foot, I have visual/tactile/auditory experiences associated with the tapping). They can also instantiate *action assessment* (my foot-tapping might be pleasant or unpleasant, seem rhythmic or not, etc.). While the presence of *freedom* ultimately depends on what freedom is, insisting that *freedom* is lacking is implausible given that (i) *mineness* is present, (ii) one does not feel forced to behave in the way one does, and (iii) one is free to stop the behaviour at any moment.

Someone may object that NCIs lack *purposiveness* and *execution*. The point about *purposiveness* can be resisted insofar as (i) allowing for a behaviour to happen can be a way of mandating it and (ii) not trying to prevent something can be a way of allowing it to happen.

While *execution* is admittedly lacking, considering it essential for action is arguably mistaken given the previously mentioned evidence for the unreliability of introspection. Since execution characterises the conscious perspective of the acting subject, the relevant sense of accomplishing a goal should be consciously accessible to the subject. And if one has a conscious sense of accomplishing a goal, one should know, and be able to reliably report, what that goal is, at least in most cases. After all, execution is 'the sense that one is *doing* what one intends to do—not merely the sense of striving towards some goal, but of successfully achieving it' (Mylopoulos & Shepherd, 2020, p. 178). However, evidence for the

unreliability of introspection shows that people very often misidentify the causes of their own decisions and actions, partially because they tend to explain their own behaviours with a priori theories about the causal connection between stimulus and response (Nisbett & Wilson, 1977, p. 233). Insofar as having a sense of accomplishing a goal involves knowing what that goal is, the evidence in question suggests that execution is not as common as it may seem. So if execution were necessary for genuine action, the latter would be a relatively rare phenomenon.

In sum, not all aspects of agentive phenomenology have to be present for a genuine action to occur, and none of them seems essential for action.

5 Conclusion

According to Phillips, many putative cases of unconscious perception face the problem of attribution. Because they do not guide action, there is no good reason to categorise them as personal states. Insofar as genuine perception is a personal state, unconscious perceptual representations do not exemplify genuine perception.

Phillips argues that unconscious perceptual representations do not guide action because their behavioural effects fail to meet certain sufficient conditions for action guidance. The behaviours in question are mere modifications of action, not self-standing actions. They are not accompanied by consciously intending (willing…) to behave in the relevant way.

I have argued that this reasoning is flawed and cannot be easily fixed. For the conditions Phillips indicates as sufficient for action guidance may not be necessary for action guidance. Indeed, when it is assumed that these conditions are necessary for action guidance, the whole reasoning behind the problem of attribution becomes unsound. Since consistency requires that Phillips applies the distinction between personal and sub-personal states/events across the board, he is compelled to apply it to action as well. And when genuine action is identified by personal-level criteria, it turns out that Phillips' conditions are unnecessary for action guidance. Therefore, even if behavioural effects of unconscious perceptual representations cannot satisfy Phillips' conditions, it does not follow that those behavioural effects do not instantiate action guidance, nor that unconscious perceptual representations are not individually attributable.

Of course, this does not show that genuine perception actually can occur outside of consciousness. But it does show that the problem of attribution is not a good reason to be sceptical about this possibility.

Acknowledgements Special thanks are due to Joshua Shepherd, the supervisor of my research stay at the University of Barcelona, where I wrote the first version of this paper. Thanks also to Benjamin Henke, Jan Kiełbasa, Matthias Michel, Ian Phillips, Assaf Weksler, the participants of the 'SemDok' seminar at Jagiellonian University, and the anonymous referees for their comments on earlier versions of this paper.

Funding This work was supported by the National Science Center (Poland) under grant number 2019/32/C/HS1/00113.

Declarations

Competing Interests The author declares no competing interests.

References

Anaya, A., & Clarke, S. (2017). Naïve realism and unconscious perception: A reply to Berger and Nanay. *Analysis, 77*(2), 267–273.

Block, N., & Phillips, I. (2017). Debate on unconscious perception. In B. Nanay (Ed.), *Current controversies in philosophy of perception* (pp. 165–192). Routledge.

Burge, T. (2010). *The origins of objectivity*. Oxford University Press.

Charles, L., Van Opstal, F., Marti, S., & Dehaene, S. (2013). Distinct brain mechanisms for conscious versus subliminal error detection. *NeuroImage, 73*, 80–94.

Dennett, D. (1969). *Content and consciousness*. Routledge & Kegan Paul.

Drayson, Z. (2012). The uses and abuses of the personal/subpersonal distinction. *Philosophical Perspectives, 26*(1), 1–18.

Endrass, T., Reuter, B., & Kathmann, N. (2007). ERP correlates of conscious error recognition: Aware and unaware errors in an antisaccade task. *European Journal of Neuroscience, 26*(6), 1714–1720.

Fiscarella, S. C., Rochet, N., & Burle, B. (2019). Becoming aware of subliminal responses: An EEG/EMG study on partial error detection and correction in humans. *Cortex, 120*, 443–456.

Frankfurt, H. G. (1988). The problem of action. In *The importance of what we care about* (pp. 69–79). Cambridge University Press.

French, C., & Phillips, I. (2023). Naïve realism, the slightest philosophy, and the slightest science. In B. McLaughlin & J. Cohen (Eds.), *Contemporary debates in the philosophy of mind* (pp. 363–383). Wiley-Blackwell.

Hursthouse, R. (1991). Arational actions. *The Journal of Philosophy, 88*(2), 57–68.

Jiang, Y. V., Costello, P., Fang, F., Huang, M., & He, S. (2006). A gender- and sexual orientation-dependent spatial attentional effect of invisible images. *Proceedings of the National Academy of Sciences, 103*(45), 17048–17052.

Johnston, M. (1997). Manifest kinds. *Journal of Philosophy, 94*(11), 564–583.

Krakauer, J. W. (2019). The intelligent reflex. *Philosophical Psychology, 32*(5), 822–830.

Logan, G. D., & Crump, M. J. (2010). Cognitive illusions of authorship reveal hierarchical error detection in skilled typists. *Science, 330*(6004), 683–686.

Mylopoulos, M., & Shepherd, J. (2020). Agentive phenomenology. In U. Kriegel (Ed.), *Oxford handbook of the philosophy of consciousness*. Oxford University Press.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review, 84*(3), 231–259.

O'Shaughnessy, B. (2008). The will. A dual aspect theory: Vol. II. Cambridge University Press.

Peters, M., Kentridge, R., Phillips, I., & Block, N. (2017). Does unconscious perception really exist? Continuing the ASSC20 debate. *Neuroscience of Consciousness, 3*(1), 1–11.

Phillips, I. (2016). Consciousness and criterion: On block's case for unconscious seeing. *Philosophy and Phenomenological Research, 93*(2), 419–451.

Phillips, I. (2018). Unconscious perception reconsidered. *Analytic Philosophy, 59*(4), 471–514.

Phillips, I. (2020). Object files and unconscious perception: A reply to Quilty-Dunn. *Analysis, 80*(2), 293–301.
Rosenthal, D. M. (2008). Consciousness and its function. *Neuropsychologia, 46*(3), 829–840.
Runyan, J. D. (2014). *Human agency and neural causes*. Palgrave MacMillan.
Schwitzgebel, E. (2008). The unreliability of naive introspection. *The Philosophical Review, 117*(2), 245–273.
Shepherd, J., & Mylopoulos, M. (2021). Unconscious perception and central coordinating agency. *Philosophical Studies, 178*(12), 3869–3893.
Watzl, S. (2017). *Structuring mind*. Oxford University Press.
Westfall, M. (2024). Constructing persons: On the personal–subpersonal distinction. *Philosophical Psychology, 37*(4), 831–860.
Wilson, T. D. (2002). *Strangers to ourselves*. Harvard University Press.
Zięba, P. J. (2019). Naïve realism about unconscious perception. *Synthese, 196*(5), 2045–2073.