

Visual Acquaintance, Action & The Explanatory Gap

Thomas Raleigh, Ruhr University Bochum

ABSTRACT: Much attention has recently been paid to the idea, which I label 'External World Acquaintance' (EWA), that the phenomenal character of perceptual experience is partially constituted by external features. One motivation for EWA which has received relatively little discussion is its alleged ability to help deal with the 'Explanatory Gap' (e.g. Fish 2008, 2009, Langsam 2011, Allen 2016). I provide a reformulation of this general line of thought, which makes clearer how and when EWA could help to explain the *specific* phenomenal nature of visual experience. In particular, I argue that by focusing on the different kinds of *perceptual actions* that are available in the case of visual spatial vs. colour perception, we get a natural explanation for why we should expect the specific nature of colour phenomenology to remain less readily intelligible than the specific nature of visual spatial phenomenology.

1. Introduction: External World Acquaintance (EWA).

Much attention has recently been paid to the following thesis, which I will label 'External World Acquaintance' (EWA):

- **EWA:** (at least some of) the phenomenal character of (some) visual experience is (partially) constituted by features located in the subject's external environment.

Whilst EWA is most obviously associated with 'naïve-realist' or 'relational' theories of perceptual experience, in fact a number of prominent representational theorists – e.g. Dretske (1996), Lycan (2001), Byrne & Tye (2006), Tye (2000, 2009, 2015), Ross (forthcoming) – have also endorsed it. These theorists claim that in virtue of having a perceptual experience that veridically represents the scene before you, the token phenomenal property-instances you are thereby aware of are, at least partially, constituted by token properties – e.g. shape and colour properties – of the perceived objects located out in the environment¹. And indeed even some sense-data theorists (Jackson 1977, Bermudez 2000) have claimed that sense-data are located out in the external environment². So endorsing EWA cuts across the standard taxonomy of theories – naïve-Realist, representational, sense-data. To be clear then, as I am using the term 'acquaintance' here I am *not* assuming that it is

¹ For example, Dretske writes:

'... the qualities that individuate one experience from another, the qualities that make *seeing* so much different from *hearing*, and seeing *red* so much different from seeing *green*, are (or need be) nowhere in the person wherein resides the experiences of these qualities. The experiences themselves are in the head (why else would closing one's eyes or stopping one's ears extinguish them?), but nothing in the head...need have the qualities that distinguish these experiences.

... the qualities by means of which we distinguish experience from one another are relational.'
(Dretske, 1996, 144-145)

² And if we think of the classic sense-data theorists, G. E. Moore famously went back and forth throughout his career as to whether or not sense-data are located out in the external environment. And though Russell in his logical atomist phase denied that sense-data are identical to (parts of) physical objects, he nevertheless wrote: "...I hold that the sense-datum is certainly something other than the subject, something to which the subject's relation is just as 'external' as to the physical object." (Russell, 1913, 78). By the time of his 'Mysticism and Logic and other Essays' (1918), Russell had become a Neutral Monist and so then identified sense-data with the ultimate physical constituents of reality.

non-representational; I just mean it to pick out that direct, intimate relation we have to elements in our own stream of consciousness, whatever the precise metaphysical nature of the relation or the stream turns out to be. EWA then is the idea that *external features*, out in the distal environment, can themselves *constitutively* enter into in the stream of consciousness and thus be objects of conscious acquaintance.

Now, although it has generated much discussion in the recent literature, I think it is fair to say that EWA remains a minority view. The orthodoxy, both in philosophy and in cognitive science, is that conscious experience supervenes solely on the brain. Much of the recent interest in EWA has been due to the prospect that *disjunctivism*, the thesis that perceptual experiences and hallucinations are fundamentally different kinds of conscious state even if they are subjectively indiscriminable, might offer a new way to resist arguments from hallucination that would otherwise push us towards orthodox phenomenal internalism³. But why should anyone want to endorse EWA in the first place? Even supposing that it could be adequately defended against arguments from hallucination and illusion, what's the positive appeal?

Many different motivations have been advanced for endorsing something like EWA. One traditional answer, in slogan form, is that something like EWA could rid us of a 'veil of appearances'. But even supposing that EWA is uniquely well placed to banish intermediaries, this just raises the further question: what's so bad about intermediaries? One suggestion is that the model of a veil of appearances is what leads us towards traditional scepticism about knowledge of the external world. And so abandoning this model of experience would give us a response to the sceptic (e.g. McDowell, 2008). For what it's worth, I'm unconvinced (see Wright 2008, Logue 2012), but I won't consider this sort of motivation any further. Nor will I consider a number of other possible motivations: EWA's alleged ability to allow us to make singular/demonstrative reference to external objects (e.g. Campbell 2002), its alleged fit with externalism about mental content (Dretske 1996, Lycan 2001, Tye 2000, 2009, 2015), its alleged ability to best explain our imaginative phenomenology (e.g. Martin 2002), or to explain our temporal phenomenology (e.g. Soteriou 2013), or the alleged shortcomings of its rivals (e.g. Travis 2004, Raleigh 2015).

A possible motivation for EWA that has received comparatively little attention is the idea that by adopting EWA we will be able to make progress with the notorious 'Explanatory Gap'⁴. The (alleged) 'Explanatory gap' concerning phenomenal consciousness is, of course, that identifying a conscious experience with some physical structure or process, perhaps in the brain, seems to leave (at least some of) the phenomenal nature of the experience unexplained. Consciousness thus provides a paradigm example of what have been called 'location problems'⁵ – phenomenal properties, like, say, moral properties, can seem peculiarly difficult to locate in the world described by the physical sciences. Fish (2008, 2009), Langsam (2011) and Allen (2016) have all claimed that a virtue of an EWA account of experience is its ability to deal with this 'Explanatory Gap' – i.e. to remove our sense of mystery or puzzlement concerning the relation of our phenomenology to the physical world. (Fish's, Langsam's and Allen's views will be discussed further in sections 4 and 6, below.)

I think that these authors are correct that EWA offers a potentially important, and underexplored, way to approach the explanatory gap. However, I will suggest that this line of thought has not yet been developed in its most persuasive form. I will argue that such theories can give us a satisfying explanation of why *spatial* visual experience has the specific phenomenal character that it has (rather than some other character), but that this form of explanation seems less satisfying for the case of *colour* phenomenology. That will lead us to the question of *why* the specific phenomenology of

³ The disjunctivist idea is due originally to Hinton (1967, 1973), though it gained more widespread discussion due to Snowdon (1981) and McDowell (1982), and receives its most sophisticated defence in the work of M. G. F. Martin (2004, 2006). See Soteriou (2016) for an excellent recent book-length survey.

⁴ The phrase 'Explanatory Gap' is due originally to Levine (1983).

⁵ The phrase 'location problem' is due originally, I believe, to Jackson (1998).

colour experience seems to remain more mysterious – i.e. to provoke a stronger sense of an unbridgeable explanatory gap – than the case of (visual) spatial phenomenology and so to resist this EWA form of explanation. And it is here I think that appealing to the nature of our visual-perceptual *sensorimotor interactions* can supply an answer as to why colour phenomenology will inevitably remain more mysterious-seeming than (visual) spatial phenomenology. Mention of sensorimotor interactions will naturally call to mind ‘Enactivist’ approaches to perception and so I will also briefly compare and contrast the position I am advocating here with the some of the claims made by some well-known Enactivists. But to begin I want to try to get this specific motivation for EWA more clearly into view and also to discuss one of EWA’s key commitments.

2. Getting the Motivation for EWA into view

Jesse Prinz (commenting on the ‘Sensorimotor’ theory of Alva Noe, 2004) writes:

“Just as it’s hard to understand why brain states feel a certain way, it’s hard to understand why brain states together with bits of the external environment feel a certain way. As Noe realises, wide supervenience will not solve the hard problem of consciousness. So why go outside the head?” (Prinz, 2006, p17)

How might one begin to answer Prinz’s challenge? What is the explanatory benefit of including parts of the external environment, as well as the subject’s brain, in the supervenience base for consciousness?

The first point to make is that the EWA model of experience does not *simply* extend the supervenience base for experience; it ascribes to experience the structure of an *interaction* between the two relata, each of which plays different roles (object and subject of experience). So it is not simply a matter of having a larger collection of things for experience to supervene on, we can also now appeal to the way or manner in which the relata *interact* (are related). Here is a very quick, initial gesture at how this *might* help. It is problematic or mysterious to simply identify phenomenal properties with neural properties (phenomenal episodes with neural processes) – how could this reddish-looking phenomenal property simply be one and the same as the pattern of firing and activation amongst these neurons? But it is equally problematic to identify phenomenal properties as simply mind-independent features of the perceived environment. There are a range of familiar problems here: perceptual relativity, the possibility of hallucination, illusion, Brains-in-Vats etc, which all congregate around the hard-to-deny fact that the state of the brain (and sensory apparatus etc) fundamentally influences, and sometimes apparently suffices on its own to produce, phenomenal experience. In short, neither kicking phenomenal properties upstairs into the physical brain nor downstairs into the mind-independent environment seems to be a satisfying strategy for locating them non-mysteriously in the physical world.

The fresh, hopeful promise of EWA is that these two options need not be exhaustive – for we can also appeal to properties of the *entire interaction* between subject and object. There need be nothing spooky or non-physical about interactions having properties. Consider a car crash – the violence of the car crash is not a property attributable individually to either car that takes part in the crash (it is not ‘located’ in either car alone); rather, it characterises the *manner* of their interaction. Likewise, the dancing of a tango can be sexy, even when neither dancer individually can be correctly so characterised. Notice also that the violence of the car crash is *intelligibly* explained as arising not only from the intrinsic natures of the two cars (e.g. their rigidity, their weight etc) but also from the way they are related to and interact with each other (e.g. moving towards each other so as to make contact at high speed). The hope then is that phenomenal properties might be plausibly ‘located’ as relational properties of the entire experiential interaction between subject and object, without implausibly identifying them as being properties intrinsic to either.

Of course this is not yet to offer any kind of solution to the ‘Explanatory Gap’. At this stage I am just indicating, in answer to Prinz’s challenge, why it *might* at least seem *promising* to try appealing to elements outside of the brain, as well as elements within, when trying to ‘locate’ phenomenology as part of the natural world. Given the famous difficulty and intractability of the ‘Hard Problem of Consciousness’, we should at least be open to exploring new ways of conceiving of phenomenal properties.

3. Dual Awareness

At this point it will be worthwhile to explicitly state an important feature of perceptual phenomenology, which I think must be a part of any plausible EWA account and which will be important later when we come to consider the differences between visual spatial phenomenology and the phenomenology of seeing colours.

The fundamental idea with EWA is that in perception we are (in some sense of ‘direct’) directly conscious of external mind-independent items and features in the environment. But, of course, EWA must still acknowledge and somehow account for the relativity of perceptual appearances. It is undeniable that the way an object looks, the way its shape or its colour looks, depends in part on features of the subject’s situation – her position, the lighting, other contextual features, focus of attention, the specific state of her visual system etc. An EWA theory then must reject any kind of inference along the following lines: since the specific way an object looks is dependent on the subject and her circumstances, therefore all the subject can really/directly be conscious of are merely the downstream causal effects of the object on the subject (her brain, her inner stream of consciousness), rather than the object itself.

It is vital then for an EWA account that the specific *way* that an external feature looks to me (from here, in this light, etc.) is not the sort of thing that could *compete* with the external feature itself to be the real, direct object of awareness. The specific phenomenal way an external feature looks/appears to me is certainly something that ‘shows up’ in the stream of consciousness, but this way of looking cannot be allowed to become an intermediary standing between the perceived feature and the conscious subject. And so, I suggest, EWA theorists must acknowledge that there is always a kind of *dual awareness* in perceptual experience: we are simultaneously conscious of the external feature and we are conscious of the specific phenomenal manner in which it is present in our consciousness. These are not simply two different objects of awareness on the same footing – rather, any instance of perceptual awareness is always, in one fell swoop, awareness of something that is being presented in some specific way/manner.

A comparison: there is no such thing as my dancing with you in no specific way or manner – as if we could just be dancing *simpliciter*. We are bound to be dancing in some specific way or manner – e.g. awkwardly or fluently, slowly or quickly, in the manner of a tango or of a salsa, etc. This token instance of the specific way or manner in which I dance with you is not something that exists independently of the actual episode of our dancing, our interacting. The way or manner in which we dance is not some further causal upshot of our interacting – like the heat we emit or the noise we make. Rather the way or manner in which we dance is constituted by our two specific natures and the specific way that we interact. And of course there is no temptation here to think that the specific way or manner in which I am dancing with you is the real or direct object that I am dancing with, and so I only count as indirectly dancing with you!

Likewise, an EWA theorist should insist that there is no such thing as conscious perceptual awareness of an external feature that presents the feature in no specific phenomenal way or manner – just pure awareness/acquaintance *simpliciter*. But equally there can be no specific phenomenal manner/way in which genuinely *perceptual* consciousness occurs without some (external) item or feature that I am aware of in that phenomenal manner/way. Thus the object of conscious awareness and the phenomenal manner in which this awareness occurs are mutually dependent and should not

be thought of as simply two distinct candidates of the same kind that could compete to be the real/direct object of awareness.

A brief, toy example here: When I look at a penny from an angle I am, in one fell swoop, consciously presented with both the circular shape of its facing surface and the specific way that that shape looks from my perspective. So the mind-independent circular shape property ‘shows up’ in my stream of consciousness; but it necessarily shows up presented in some specific phenomenal manner – a relational feature of the experiential situation. It is still the penny’s intrinsic circular shape that I am (directly) aware of when it is presented to my mind in this specific phenomenal way of looking.

Now of course, theorists who hold that experience supervenes solely on the brain can allow that the way an object looks/appears is the *causal* upshot of both the nature of the object and the nature of the subjective viewing conditions. But it is a distinctive commitment of EWA that the specific way that the object looks (from here, in this light etc) is *constitutively* dependent on both object and subject – just as the manner in which two dancers are dancing constitutively depends on the nature and movements of both dancers, or the manner in which two cars crash constitutively depends on the nature and movements of both cars.

Of course, theorists sympathetic to EWA are well aware that they must account for perceptual relativity and for the joint contributions to phenomenal character of both the external objects of perception and the manner in which they are perceived⁶. Nevertheless, I don’t think that this dual-aspect nature of perceptual consciousness has always been sufficiently clearly recognised as an essential commitment of EWA theories. In the very same instance of conscious perceptual awareness one is presented both with an external, mind-independent feature and with the specific phenomenal way that it looks/appears to you (given that viewpoint and sensory modality and pattern of attention, etc.), which is a mind-dependent, relational feature of the whole perceptual situation.

By claiming that this kind of dual-awareness is an essential part of any plausible EWA theory, I am thereby advising EWA theorists to reject a particular *strong version* of the ‘Transparency of Experience’ thesis, according to which the *only* phenomenal features of a perceptual experience that one can ever gain introspective awareness of are the external objects and properties (etc.) that one is (apparently) seeing. Notice however that this kind of dual-awareness is consistent with a weaker understanding of Transparency, which holds only that whenever one introspectively attends to a phenomenal feature in perceptual experience one cannot do so without thereby attending to some (apparently) external feature that is being perceived⁷. (One cannot attend to the specific phenomenal manner in which one is aware of O without thereby attending to O itself.)

⁶ Matthew Soteriou comments:

‘...there may be different ways in which the conscious characters of one’s conscious sensory experiences may vary—some differences may be due to differences in the sorts of entities one is psychologically related to, and some differences may be due to the different ways in which one is psychologically related to them.’ (Soteriou, 2013, 114)

In the same vein, John Campbell (2002, 2007, 2009) has provided an influential account that appeals to a 3-place relation of perceptual acquaintance, relating subject, object and a viewpoint. See also Kennedy (2007), Fish (2009) and Brewer (2011), all of whom acknowledge (in different ways) that such factors as: the state of the subject’s visual system, the subject’s attention, the subject’s conceptual dispositions, can contribute to perceptual phenomenology.

⁷ Soteriou (2013) expresses the distinction between the stronger and weaker versions of the transparency thesis thus: ‘[Transparency] is sometimes expressed in terms of the claim that introspection of one’s perceptual experience seems to one to reveal *only* the objects, qualities, and relations one is apparently perceptually aware of in having that experience, and not qualities of the experience itself. A weaker version of the claim is that when one attempts to attend introspectively to what it is like for one to be having a perceptual experience, it seems to one as though one can only do so through attending to the sorts of objects, qualities, and relations one is apparently *perceptually* aware of in having the experience.’ (Soteriou, 2013, 88)

Rejecting the stronger version of Transparency is, I suggest, a result to be welcomed since the thesis is anyway implausible. For it seems clear, in the case of visual experience, that the perspectival, spatial nature of the visual field is a phenomenal feature of experience that we can be introspectively aware of. Matthew Soteriou puts this point in terms of the boundary or limits of the visual field:

‘The fact that the spatial sensory field of vision is a phenomenological feature of visual experience gives us reason to reject the stronger version of the [Transparency] claim. The boundaries of this spatial sensory field are not features of some object one is visually aware of in having a visual experience. One cannot directly attend to them in the way in which one can directly attend to the objects and features that fall within them. But they do nonetheless feature in the conscious character of one’s visual experience, and one can become aware of them when attending to the objects that fall within them.’ (Soteriou, 2013, 119)

4. The Basic Strategy For Closing The Gap

Bill Fish (2008, 2009) suggests we can think of the explanatory gap as coming in two parts:

- Why is it like *anything*? Why is it like *something or other* to be in this brain state/functional state? Why is there a conscious subjective perspective at all? Why am I not a zombie?
- Why is it *specifically* like *this*? Why does it feel/appear *this specific way* when in this brain/functional state, rather than some other specific way? Why do I not have, say, an inverted spectrum?

Fish cites a number of theorists – Shoemaker (1975), Tye (2006), Clark (2000), Kirk (2005), Kim (2007) – who claim to give a broadly functionalist solution to the first problem. All of these philosophers are essentially offering arguments against the possibility of ‘zombies’ – functional duplicates of a conscious subject who nevertheless lack consciousness altogether. However at least three of these philosophers (Tye, Clark and Kim) concede that whilst the functionalist story rules out zombies, it does not totally solve the hard-problem for it does not solve the ‘Why is it *specifically like this*?’ problem – i.e. it does not rule out functional duplicates with ‘switched around’ phenomenology, as in ‘Inverted Spectrum’ cases.

I will not make any attempt here to evaluate or defend this family of functionalist arguments against zombies/epiphenomenalism. How does EWA come into the picture? Fish suggests that *if* these kinds of functionalist arguments not only provide a satisfying explanation of why we enjoy *some kind* of subjective conscious perspective, but can also be adapted so as accommodate EWA, we could then get the missing answer to the ‘Why is it *specifically like this*?’ problem. For we could now appeal to the extra, environmental element in experience to help determine the *specific* phenomenal nature of the experience. The reason that being in this neural/functional state feels/appears *some or other way*, rather than no way at all, is that the functional state guarantees that I am consciously acquainted with something or other. (I am enjoying a conscious experience that constitutively involves some feature in my external environment.) But the reason that my experience feels/appears *this specific way*, rather than some other specific way, is that I am acquainted with *this specific environmental feature*, rather than some other external feature. So the explanation for the specific phenomenal nature of experience appeals to more than just the functional state of the subject, which only generates/guarantees conscious receptivity to something or other; it also appeals to the nature of the perceived object – a further, constitutive element of the experience, according to EWA. Fish sums this line of thought up as follows:

“The reason that it is like that to see green is because it is real-world greenness that the system’s functional processing serves to acquaint it with – we no longer need to see the functionalist programme as attempting to conjure phenomenal greenishness out of function. This also explains why God had no more work to do once He had fixed all the physical and functional facts. In fixing these facts He both made it the case that certain creatures in the

world would be acquainted with elements of their environment and fixed the nature of the properties that they are acquainted with.” (Fish, 2008, 174-5)

Likewise in his book, Fish writes:

‘The difference in what it is like to see a ripe McIntosh apple and what it is like to see a ripe cucumber is not explained by the differences in [neural/functional] processing – instead, it is explained by the different color properties that the two objects possess. When we see a ripe McIntosh apple, the phenomenal character of our experience is its property of acquainting us with the fact of the object’s being *red*; when we see a ripe cucumber, it is the experience’s property of acquainting us with the fact of the object’s being *green*. This is where the difference between what it is like to have the two experiences is to be found.’ (Fish, 2009, 75-6)

‘This approach therefore offers to provide the most intelligible, transparent explanation yet of just how it could be that the truths about phenomenal consciousness supervene on the physical/functional.’ (ibid. 78)

Harold Langsam (2011) makes a number of similar claims on behalf of the relational theory of experience:

“...it is only the combination of an observable property with consciousness that results in a phenomenal property, a phenomenal property that is intelligibly similar to the observable property from which it is derived.” (Langsam, 2011, 54)

“...what it is like to be conscious of redness, is intelligibly explained by the combination of consciousness and redness in the structure of an experience of red.”. (ibid, 55-56)

“reflection on the natures of redness, phenomenal redness, and consciousness reveals that the nature of phenomenal redness is such as to result in an intelligible way from the combination of an act of consciousness directed at redness. The result is not only an intelligible one but a necessary one. Given the natures of redness and consciousness, it is necessary that their combination result in the instantiation of the phenomenal property we have been referring to as phenomenal redness.” (ibid, 59)

And more recently, Keith Allen (2016) also claims that naïve-realism can help close the explanatory gap:

“According to the naïve realist theory of perception, colour experiences are essentially relational events that ‘inherit’ their qualitative character from the qualitative character of the colours that they are experiences of. As such, what it is like to perceive colour is determined, quite simply, by the nature of the colours perceived.

... this account takes seriously what Chalmers calls the ‘hard problem of consciousness... According to the naïve realist, the hardness of the hard problem derives from the fact that neural processing is not itself sufficient for perceptual experience, but merely an enabling condition of it: perceptual experiences have the qualitative characters they do because of the objects and properties in our environment that they acquaint us with.” (Allen 2016, 181-182)

To briefly recap: the general strategy here is to take the existence of a conscious relation of awareness or acquaintance for granted – i.e. take for granted that we have an answer to the ‘why is it like anything?’ question – and then seek to explain the specific phenomenal nature of specific experiences by appeal to the specific external features that one is consciously acquainted with. To the extent that the qualities one is acquainted with are relatively non-mysterious, we thereby gain some explanatory traction with the ‘why is it specifically like this?’ question.

5. Does this kind of strategy really require EWA?

Before continuing, we ought to consider a question that Fish (2008) attributes to Heather Logue: why couldn't a standard – i.e. phenomenally *internalist* – representational theorist make a parallel appeal to the worldly features that are *represented* in experience – and thus give just as good an answer to the specific 'what it's like' question, without committing to anything like EWA? I think that Logue (2012) herself, who in turn cites Mark Johnston (2006), gives the right sort of response:

“...as long as phenomenal character consists in representations of things, we can't rule out that what it's like to perceive something is radically different from what that thing is like in itself. ... For representations can of course be totally different from what they represent. So we couldn't be sure that the phenomenal character determined by a representation of crescent-shapedness provides any guide to what crescent-shaped things are like in themselves. As Johnston would put it, there would be no way of ruling out the possibility that the human way of representing crescent-shapedness, and the phenomenal character it plays a role in determining, is 'entirely idiosyncratic relative to the intrinsic natures of things'.” (Logue, 2012, 13)

In other words, if our task is to explain why a conscious experience has the specific phenomenology it does, we cannot just appeal to the fact that the experience represents F-ness, for this fact provides *too little constraint* on phenomenology. In general there will be indefinitely many possible vehicles, mediums or manners of representation for some bit of content – the nature of the representation is, at best, constrained only by a very weak requirement for some kind of homomorphism with what is represented. Whereas, appealing to the experience *acquainting* us with F-ness imposes a much more determinate constraint on phenomenology, as this actual instance of F-ness is supposed to partially constitute and 'show-up' in the phenomenology.

Now at this point, a certain kind of representational theorist – a so-called 'strong intentionalist' – might want to protest along the following lines: there may indeed be indefinitely many different possible vehicles or mediums for representing some specific content, but in having a perceptual experience with a representational content, you are not thereby aware of the *vehicle* or *medium* of representation – rather, you are simply aware of the representational *content*. In other words, your visual phenomenology is exhaustively constituted/determined by the visual experience's representational content *that* the environment is some specific way. Notice that the thesis of strong intentionalism, that phenomenology is wholly determined by content, is *different* from the representational variety of the EWA thesis mentioned back at the start – which holds that in a veridical perception, in virtue of accurately representing the scene before you, the token phenomenal properties presented in your experience are constituted by token external properties (shapes and colours etc.) of that environmental scene. Whilst these two ideas are compatible, a representational theorist can certainly endorse strong intentionalism without endorsing EWA – for one might hold that phenomenal properties supervene entirely on representational content without accepting that those phenomenal properties are constituted by token instances of properties located out in the environment.

And thus such *non-EWA*, intentionalist theorists might feel that they are equally well placed to account for the specific phenomenal nature of a perceptual experience. Your experience has the specific phenomenology it does because it has a specific representational content that the environment is some specific way – and this content *exhaustively* determines the phenomenology.

In response I would make three points:

(i) Firstly, the idea that one can be conscious only of pure representational content without being able (even if you strive and concentrate!) to gain any conscious awareness at all of the vehicle or medium or manner of representation (i.e. any of its non-representational properties) is hardly a

straightforward or immediately intelligible idea. We have no everyday familiar model for such a thing. In everyday examples of physical representations – paintings, photos, written sentences, spoken words, sign-language, sculptures, maps, smoke signals, traffic lights etc – some intrinsically non-representational physical property – e.g. shape and/or colour properties of the vehicle/medium, or acoustic and temporal properties in the case of spoken words – are *used* to represent (stand for) something. In these real-life cases we gain awareness of representational content *in virtue of* being conscious of some of its intrinsically non-representational features. Now, appealing to these sorts of cases does not, of course, demonstrate that it is flat-out metaphysically impossible for a subject to have ‘pure’ awareness of a representation’s content without having any awareness of the representational vehicle’s intrinsic/non-representational features. But nor should it just be granted that appealing to this sort of awareness of pure content, divorced entirely from the concrete vehicle or method of representation, is an obviously legitimate or persuasive move in the dialectic.

(ii) Secondly, the idea of being directly phenomenally conscious of pure content – i.e. *not in virtue of* being conscious of any aspect of the concrete vehicle/medium for this content – is also a *metaphysically* puzzling idea. For contents are, on virtually any theory, *abstracta*. They are propositions, or sets of possible worlds. Even a non-propositional content – such as a so-called ‘gappy content’ – would presumably also be something abstract that could be shared/instantiated by a whole range of different distinct concrete mental states or representations. And bearing a relation to something abstract simply doesn’t look like the kind of thing that *could* explain the vivid real specific nature of conscious phenomenal character. David Papineau (2014) writes:

“What is it for a subject to bear some sensing relation to a proposition, or alternatively for some sensory vehicle state to express a proposition? ... propositions are abstract entities. There seems something quite amiss with the suggestion that my here-and-now conscious feelings are constituted by my bearing any kind of relation to abstract entities. ... My conscious sensory feelings are concrete, here-and-now, replete with causes and effects. How can their metaphysical nature essentially involve relations to entities that lie outside space and time?” (Papineau, 2014, 6-7)

(iii) Thirdly, we can consciously think/judge that the environment before us is some specific way. But whatever cognitive phenomenology (if any) such an episode might have, it is surely phenomenally *very* different from visually experiencing the environment being that specific way. There is thus a strong *prima facie* case against the idea that a conscious state’s representational content by itself can exhaustively explain its specific phenomenal nature. For whatever our visual experience represents about the environment, it is plausibly possible to think/judge that the environment is that way – yet the phenomenology of visual experience and of judgement are utterly different.

6. Improving the strategy

I think that the basic EWA strategy, outlined in section 4, contains an important insight as to how the specific nature of phenomenology could be *intelligibly explained* as the joint upshot of the object and the conscious subject. However, by choosing *colour* as their example of an environmental property that we can be consciously acquainted with, I think that Fish, Langsam and Allen have all launched this general EWA strategy from an unpersuasive base.

Langsam says almost nothing about the metaphysics of colour⁸, though it is clear he assumes that there are such familiar, observable properties of external objects. Fish briefly discusses the metaphysics of colour and makes clear that as a naïve realist he is committed to a realist/physicalist

⁸ At one point Langsam does claim, in passing, that the manifest sensory nature of a colour is distinct from the ‘the microstructural properties that underlie them’ (Langsam, 2011, 31). This might be interpreted as evincing a kind of primitivism about colours. See the discussion of primitivism and of Keith Allen’s work later in this section.

treatment of colours as ‘aspects of the objective physical world’⁹. But there is a range of familiar problems with treating colour as a mind-independent quality of external, physical objects. To rehearse just two: it is not at all clear that all the items which look green to us (in normal or optimal conditions of viewing) have any intrinsic, natural (non-disjunctive) property in common. This is the problem of metamers – very differently textured surfaces can reflect quite different combinations of wavelengths of light and yet give rise to identical (indiscriminable) colour experiences/appearances. (A related issue: colours are not only seen via light *reflected* from surfaces, we also see colours from light-emitting objects – e.g. coloured light bulbs – and from light-refracting or light-diffracting volumes – e.g. a glass of red wine, a volume of dust-filled air.) Secondly, there seems to be nothing intrinsic to physical light or in the physical surfaces (or volumes etc) that accounts for the divisions and relations between the colours as experienced – i.e. for the 4+2 structure of the primary colours. We experience all colours as mixtures of the 4 primary colours – red, blue, green, yellow, none of which can be further analysed as mixtures of colours – plus the 2 non-spectral factors, black and white. The 4 primary colours have pretty specific locations amongst the wavelengths of light – most subjects agree to a high degree of accuracy that blue is at about 465 nanometres, green at 505 etc. But there is nothing *intrinsically* special about these wavelengths of light! The spectrum of visible light, after all, is a continuous range of electro-magnetic energy – there are no natural divisions or structure in the spectrum itself. It seems then that any explanation for the 4+2 structure would have to appeal to features of the human visual system – as a first approximation, to the ‘opponent-process’ mechanisms in the eye¹⁰.

These sorts of considerations are commonly wielded against physicalism/realism about colour, but it is not my intention to take any positive stance here about the metaphysics of colour – I certainly do not wish to suggest that the above-mentioned issues conclusively rule out any form of physicalism/realism about colour. The point is just that an advocate of the EWA-strategy with the Explanatory Gap cannot just *assume* that we have a clear, unproblematic grip on what the mind-independent quality (if indeed there is such a quality) of environmental colour is, *apart* from how it appears in our colour experiences. The large literature and widespread disagreement about the metaphysics of colour testifies to this. And so appealing to such an (alleged) quality in order to remove the mystery surrounding our specific colour phenomenology is liable to leave us with no less of a sense of mystery.

We are told by Fish and by Langsam that a colour experience is *this* specific way as we are acquainted with a specific environmental property, but when we try to say *which* environmental property it is that we are acquainted with in, say, experiences of green, it is at least highly problematic to find anything that will do the job. The problem here is twofold. Firstly, we end up facing a location problem with external colour qualities that is not obviously much less daunting than the location problem concerning phenomenal properties with which we started¹¹. Secondly, even if a convincing case could be made for identifying environmental colour with, say, some kind of micro-textural, light-reflecting property of surfaces, it is not at all obvious how/why acquaintance with *this* specific kind of property would remove our puzzlement about the specific nature of our colour phenomenology. For it seems we could still legitimately ask the specific What-Its-Like Question –

⁹ See Fish, 2009, 150-159. Fish ends up tentatively endorsing the idea that in addition to an intrinsic physical colour, objects also have ‘shades’, which are meant to be relational but still mind-independent properties that can vary as the surrounding lighting conditions vary.

¹⁰ Research on the precise details of opponent-process theory is ongoing and it seems likely that other ‘post-receptor’ processes will also be relevant – e.g. lateral inhibition mechanisms may explain some colour contrast effects (Ratliff, 1965). Nevertheless, that opponent-processes will be one very large and important part of the story is widely accepted as beyond dispute in vision science. For an introductory account of opponent processing see Ch.3 of Palmer (1999). For some remaining questions about how exactly opponent processing works see Byrne & Hilbert (2003).

¹¹ See Byrne (2006), for a fable about a possible world where philosophers debate the ‘colour-body problem’ in a way that exactly parallels our actual world debates about the ‘mind-body’ problem.

why should acquaintance with, say, *this* specific surface-texture/reflectance property visually appear *this way* rather than some other specific way?

Now, microscopic reflectance properties are not the only possible candidate properties for a colour physicalist. One might instead be a ‘primitivist’ about colour and resist identifying it with, or reducing it to, any other kind of physical property at all, maintaining instead that colours are *sui generis*, simple, irreducible physical properties of objects – see e.g. Hacker (1987), Campbell (1994, 2005), McGinn (1996), Gert (2008). And indeed this is the theory that Allen ably defends in his 2016 book, where he also advances the idea that if such a primitivist view of colour is combined with a naïve-realist account of perceptual experience, we get a nice explanation for the specific phenomenal character of colour experience. Again: I want to remain, as far as possible, neutral about the metaphysics of colour so it is not my intention here to provide any fresh arguments against primitivism. And it should certainly be admitted that *if* one had a convincing argument, *independent* of trying to defend any kind of EWA or naïve-realism about experience, for accepting that there are such primitive, *sui generis* colour properties out there in the world, then it might be legitimate to appeal to these primitive colour properties as part of a strategy for dealing with the explanatory gap.

However, primitivism remains very much a minority view so it would be a significant cost for the present EWA strategy if it required one to be a primitivist about colours. More importantly, whatever the general pros and cons of primitivism as a metaphysical theory of colour, the problem with using it as part of the present EWA strategy with the explanatory gap is that we risk ending up with a ‘*virtus dormitiva*’ kind of explanation. The idea was to understand the mysterious phenomenology of colour experience by appeal to some relatively *non-mysterious*, mind-independent feature that we are acquainted with. But now the candidate external feature turns out to be a special primitive, *sui generis* property that just happens, as a matter of brute metaphysical fact¹², to supervene on the rest of the object’s physical properties. Moreover, these primitive colour properties don’t really play any other theoretical role, nor can they be given any further substantial characterisation, *except* as being those features that we perceive in our colour experiences. For anyone not already convinced of the virtues of colour primitivism, this will surely feel like an uncomfortably tight explanatory loop. Bear in mind also that Allen’s brand of primitivism (following Kalderon, 2007) does not just ascribe a coloured surface *one* primitive colour property; rather, a coloured surface is supposed to simultaneously possess a whole *range* of such primitive properties to account for the variations in colour appearances between different observers in different conditions, with each different perceiver’s visual system ‘selecting’ just one from these many different primitive colour properties on each occasion. In terms of ontological ‘expense’, positing a multitude of extra, primitive colour properties, whose supervenience on other physical properties cannot be further explained, does not seem to be much less extravagant than the externalist varieties of the sense-data theory (Jackson 1977, Bermudez 2000), which claim that colourful sense-data are somehow located out on the surfaces of physical objects, whilst remaining ontologically distinct from them. I suspect that most philosophers who are interested in dealing with the Explanatory Gap, whilst keeping faith with some kind of naturalism or physicalism, will feel that allowing such a multiplicity of extra primitive properties into our ontology, like allowing for externally located sense-data, is just too much to swallow.

I propose that this EWA strategy for dealing with the Explanatory Gap can be started out on a much firmer footing if our initial, parade-case example is the experiencing of *shape* (spatial properties) rather than colour.

¹² In response to the question of ‘why physical properties of objects realize the colours that they do’, Allen denies that the supervenience relation in question requires or admits of any further explanation and suggests it is something we just have to accept with ‘natural piety’ – see Allen, 2016, 182-183.

Firstly, shape is not *proprietary* to vision as colour is. We can feel a shape as well as seeing it¹³. (In scholastic terms shape is a ‘common sensible’ whereas colour is a ‘proper sensible’.) So we have some *independent* grip on what the nature of a shape is and how it might ‘show up’ in the stream of consciousness, quite apart from whatever we can glean from seeing shapes and experiencing their visual appearances.

Secondly, there is much less metaphysical puzzlement with ‘locating’ shape as a mind-independent quality out there in the external physical world, compared either with locating colour or with locating phenomenal properties as part of the physical world. So the intended explanans (a physical object’s shape) is, as it should be, significantly less mysterious than the intended explanandum (visual shape phenomenology). Now at this point I can imagine a worry as to whether the penny’s macroscopic circular shape really is any less mysterious than its colour. E.g. one might worry that the sort of macroscopic shape we perceive the penny as having is a smooth, solid, gapless disc with precise, sharply defined edges and surfaces. Yet science tells us that the penny is really composed of a scattered cloud of micro-particles with spooky quantum level properties. And so you might think that the sort of apparently perceptible macroscopic shape qualities that I am appealing to in order to explain the specific nature of our spatial phenomenology are not so non-mysteriously part of the external world after-all. Limitations of space do not permit nearly as full a discussion of this issue as it deserves. So all I will say here in response is that I think we should heed Putnam’s (1975) well-known advice that there need be nothing unscientific or bogus about explanations that advert to ‘higher-level’ macroscopic shape properties – sometimes (i.e. for some explanatory purposes) the best explanation for why a square peg won’t fit through a round hole does *not* refer to micro-particles, but simply to the geometry of the familiar, macroscopic shape properties involved. Whatever interesting, tricky issues there are concerning the relationship between the macroscopic shape of the table and the swarm of micro-particles ‘arranged table-wise’, this macroscopic shape does not give rise to anything like the same metaphysical ‘location problem’ that arises with the colour of the table. And so I think there is no problem with treating the sort of familiar macroscopic spatial features that we perceive as being at least *relatively* non-mysterious, unproblematic features of the external world whilst we are engaged in the project of trying to explain the mysterious phenomenal properties of experience.

Thirdly, and most importantly, there is no intuitive sense of mystery over the *relativity* of shape appearances. We have long had a general and precise theory specifying how shape appearances depend on both the intrinsic shape that is seen and the positioning of the seeing subject – viz, the theory of perspective. But it is not as if one has to read a treatise on perspective in order to have an understanding of the relativity of shape appearances. We all have a *natural, implicit* and *practical* understanding of how the appearance of everyday shapes and objects depends on the angle and distance of viewing. It is familiar and un-mysterious to us that a circular penny, viewed from a certain angle ‘looks elliptical’ – i.e. looks the way that an ellipse looks when viewed head-on. I shall return to the ‘familiarity’ of perceptual shape relativity in section 7, below.

The general EWA strategy with the explanatory gap thus seems to work much more satisfactorily with shape appearances than with colour appearance.

- Q: Why is my experience of the shape of the penny specifically like *this*?
- A: Because you are acquainted with a circular-shaped surface, viewed from an angle of 45 degrees and so the circular shape projects an elliptical profile towards your viewpoint.

There is, I submit, no remaining sense of mystery or explanatory gap here concerning the *specific* What-It’s-Like question. (Recall: we’re assuming – and it’s a large assumption! – that we can give a suitable functionalist account of why we have any kind of conscious perspective in the first place.)

¹³

And perhaps we can also *hear* changes in spatial location.

But compare:

- Q: Why is my experience of the colour of the ripe tomato specifically like *this*?
- A: Because you are acquainted with a physical surface with a certain microscopic texture that reflects light with a wavelength of 650 nanometres towards your viewpoint.

But . . . why should having a conscious perspective on a surface with *that* specific kind of micro-texture look *this specific way*? Our puzzlement here remains unresolved.

I think that this difference in relative mysteriousness between shape and colour phenomenology is actually an important feature of the (apparent) ‘Explanatory Gap’¹⁴. Indeed it should count as a point in favour of a theory of conscious experience if it can provide a natural explanation for this difference in mysteriousness. (Such an explanation will be offered in section 7, below.)

Why can’t other theorists appeal to the theory of perspective in the same way? After all, adherents of EWA have no special proprietary claim on the theory of perspective. But only EWA can help itself to the external instance of circularity, and the relational property of *its* projected profile relative to a viewpoint, as being *constitutive* of phenomenology. All theories which want to identify the phenomenal property (way of looking) with some state or feature ‘located in’ the perceiving subject face the question: but why should the brain state, or representational state, that is the downstream *causal effect* of an elliptical pattern of stimulation at the retina have a phenomenology that is specifically like *this*? To repeat: neither neural nor representational properties of the subject’s mental state seem able to satisfactorily answer this question. Only EWA, by invoking a relation of acquaintance between conscious subject and external object can answer the why-specifically-like-*this* question simply by appealing to the familiar shape property plus the theory of perspective *and stop there!* All other theories have to tell some further story about how an internal state of the subject, which is a downstream causal effect of the object and the viewpoint, is what determines/constitutes the phenomenology, and it is precisely this further part of the story that looks like it is bound to be unsatisfactory when it comes to explaining the specific phenomenal nature of the spatial experience.

7. Action and the difference between spatial and colour phenomenology

But so what *should* an adherent of EWA say about colour experience? And how is the account different to the EWA account of (visual) spatial experience?

Let’s begin by sketching the general metaphysical story that an EWA-theorist should tell about colour perception. Recalling the general structure of dual awareness described in section (3), a visual perceptual experience of colour (directly) presents an external, non-relational feature of some object – e.g. some kind of micro-structural/reflectance feature of the texture of a surface in virtue of which it selectively reflects only some wavelengths of light. This is the external feature that one gains conscious awareness of in a perception of (surface) colour. But the experience is bound to present this external feature in some specific phenomenal manner – i.e. looking some specific way – due to the specific nature of the perceptual interaction – i.e. the specific nature and state of the perceiver’s visual system, the specifics of the lighting, other contextual features, maybe also attention, etc. This specific phenomenal way of looking is a relational feature.

Our everyday use of colour terms sometimes seems to refer to an intrinsic physical feature of objects – e.g. “My hair is turning grey, but thankfully nobody can see it in this light”. But they can also sometimes seem to be referring to a relational feature – e.g. “The roof of the chapel turns a delicate

¹⁴ Likewise, it is commonplace to note that Jackson’s (1982) famous thought experiment about ‘Mary’, would not work nearly so well if it concerned a missing shape experience rather than a missing colour experience.

shade of pink at sunset”. As Tye (2000) and Fish (2009) both note, this is analogous to how everyday talk of ‘weight’ is sometimes best understood as referring to the object’s mass (a non-relational property), but other times it refers to weight properly so-called (a force that is the joint upshot of the object’s mass and gravity). When it comes to the metaphysical question ‘but is colour *really* a physical property intrinsic to an object or is it *really* a relational/mind-dependent feature?’ it is an advantage of the present account that it can *remain neutral*. Given the structure of dual awareness that EWA is committed to, in a perceptual experience of a coloured opaque surface we have awareness *both* of a non-relational feature of the physical surface *and* the relational property of how this surface feature looks (to a human with this kind of visual system, in these lighting conditions, within this surrounding context etc.) Thus, the foregoing account could be acceptable to colour physicalists – who wish to identify colour with the former kind of physical-surface property¹⁵. But it could also be acceptable to dispositional or ‘relational’ theorists about colour – who hold that colour is an essentially relational property, which must be defined relative to some class of perceivers¹⁶.

The EWA story for colour perception then is, at a general metaphysical level, the same as for (visual) spatial experience. There is acquaintance with a specific external feature in some specific circumstances and there is dual awareness both of that external feature and of the resulting specific way that the external feature looks under those specific conditions of viewing. There is thus, on this EWA story, no deep metaphysical difference between seeing shapes and seeing colours, such as treating the former as a primary and the latter a secondary quality. We have already had occasion to mention metamers – intrinsically quite different environmental features that nevertheless give rise to the same kind of colour experience. Many theorists have concluded that the only thing unifying these disparate environmental features is their common effect on the human visual system. But shape perception is not so very different in this regard. It is, after all, a well-established optical fact that indefinitely many differently shaped objects (or arrays of objects), positioned at different angles and distances from the subject, can end up producing the same pattern of stimulus at the retina. So intrinsically very different environmental shapes can give rise to similar shape phenomenology due to their common effect on the visual system. (This is a simplification of course – much ingenuity would be needed to counteract all the various other cues the visual system uses such as shading and colour contrast etc.)

But so then, recalling the pair of questions and answers at the end of the previous section: why is it that in the case of colour experience appealing to acquaintance with some specific external property (in a specific manner) does not remove our sense of an explanatory gap as it does with (visual) spatial experience?

Firstly, to repeat, physical shape is a familiar, relatively unmysterious macroscopic property that we can consciously experience in more than one sense-modality. And so we are able to ‘triangulate’ in, as it were, on the nature of the intrinsic property via these different kinds of shape experiences, which in turn provides us with a kind of sensory grasp of the relation between the intrinsic shape property and the experienced shape appearances.

Secondly, the major sort of contribution to visual shape phenomenology that the *subject* makes is also something familiar and macroscopic – viz. the relative position of the viewer’s eyes with respect to the object. Notice: this is also something we can *consciously experience* via more than one modality. I can reach out and feel the distance/angle to an object as well as seeing the distance/angle. I can consciously experience the movement and positioning of my eyes and head via kinaesthetic and proprioceptive sensing.

But thirdly, and most importantly, both the object’s contribution and the subject’s contribution are factors that can be *smoothly varied at will by the subject*. We can (and constantly do) vary the angle

¹⁵ An early *locus classicus* for physicalism about color is Thomas Reid (1822/1970). For more recent defences of the view, see e.g. Armstrong 1969, Tye 2000, Byrne & Hilbert 2003.

¹⁶ For an excellent recent book-length defence of a relational theory of color, see Cohen (2009).

and distance from which we perceive an object's shape. I can easily turn the round penny to any angle I like. I can walk around the square table at my leisure, bend down, step back, tilt my head etc. Thus the major contribution made by the state of the subject's visual system to the experienced shape phenomenology is readily comprehensible, for we have (in general) complete voluntary control over the crucial factor – the relative position of the visual system. And it is also, in general, relatively easy and utterly familiar to alter an object's macroscopic shape by bending or stretching or pulling or cutting etc, whilst keeping one's viewpoint fixed.

In other words, the *way* or *manner* in which the object's shape and the perceiving subject's viewpoint *combine* in an experience, so as to give rise to *this* specific phenomenology, is non-mysterious and intelligible because both the object's shape and the visual system's position are familiar macroscopic features that we have independent conscious access to and that it is (in general) easy for us to *vary independently and at will* whilst consciously experiencing the results. That is why we all have an implicit, practical and *sensory* grasp of the relationship – a relationship that is explicitly described by the theory of perspective – between the object's intrinsic shape, one's viewing position and the resulting shape appearances. For we are constantly making small, smooth adjustments to one or other factor and observing the phenomenal upshot throughout our everyday lives.

In contrast to macroscopic shape and spatial properties, the standard sorts of candidates for the mind-independent feature that we perceive when we perceive colour are *microscopic* features of the object's surfaces¹⁷ – not something familiar and macroscopic. Moreover, these microscopic textural/reflectance features are not something that we can consciously experience via any sense other than vision. But most important is the fact that within the contribution to colour phenomenology made by the subject, the main/crucial factor is something that the subject has *no independent conscious access to* and has only *very limited control over or ability to vary*. The sorts of opponent processing cells that explain the divisions and relations between the experienced colours are entirely unfamiliar and microscopic – for they take place inside the subject's brain. We cannot, as we can with shape appearances, hold the physical surface property constant and smoothly manipulate our colour processing mechanisms along various dimensions¹⁸! Nor do we have *independent* conscious access to the state of our colour processing mechanisms via some other sense modality, as we do have independent conscious access to the position of our head. And whilst we can to some extent manipulate the micro-structural properties of surfaces that contribute to colour experience in various ways, in comparison to our manipulations of macroscopic shape these manipulations of colour are generally performed indirectly, with a degree of blind, trial and error about them. For example, we may come to learn that moistening this kind of surface makes it darker in colour, that leaving this kind of surface in the sun makes it lighter in colour, that stretching this kind of surface changes its colour, etc. But we don't have any intuitive grasp of *how* or *why* moisture, sunlight or stretching etc. should have any effect on the colour – we have no idea what sort of effect these interventions are having on the micro-texture of the surface. Thus the way in which *this* kind of micro-textural property interacts with *this* kind of perceptual apparatus so as to give rise to specifically *this* kind of colour phenomenology remains mysterious and seemingly arbitrary.

¹⁷ Most current versions of colour-physicalism (e.g. Byrne & Hilbert 2003, Tye 2000) hold that colours are light-reflectance properties – i.e. a *functional* property that can be multiply realized by different kinds of micro-physical textures. Still, in perceiving such a functional property, we would be perceiving something that in any particular case is *realized* by textures of the surface at the microscopic scale.

¹⁸ Of course, *some* aspects of the subject's contribution to colour phenomenology are, as with shape, relatively everyday and familiar. E.g. We can squint or narrow our eyes so as to receive less stimulation and so things 'look darker'. But most of the important factors involved in the subject's contribution – e.g. the 3 different types of cones in the retina, opponent processes in the visual system – are not like this.

Let me now try to clarify what exactly is being claimed in the foregoing account of the difference between spatial and colour experience.

The key claim is not that we lack *theoretical* knowledge of the facts about colour perception, which we have in the case of spatial perception (e.g. the theory of perspective). The point is that in the case of spatial perception we have both: (a) independent *conscious access* to both ‘sides’ of the interaction – macroscopic shape of the object, relative position of the head, and (b) practical know-how and ability concerning how to perform a wide range of manipulations on both sides of the perceptual interaction *whilst it is being consciously experienced*. Thus we should not expect our sense of arbitrariness about the specific nature of colour phenomenology to disappear just by gaining ever more theoretical, scientific knowledge about the processes involved in colour vision. On the account being presented here, we should only expect our sense of an explanatory gap concerning the specific nature of colour phenomenology to disappear if (*per impossibile*) we could somehow independently vary and manipulate both our colour processing mechanisms and the microscopic textures of surfaces whilst consciously perceiving colour in the same easy and wide-ranging way that we can vary our perspective and manipulate macroscopic shape while enjoying a spatial experience. (And perhaps we would also need the sort of independent conscious access, via different sensory modalities, to both sides of the interaction that we have in the case of spatial experience.)

It is also worth clarifying that, whereas the basic EWA strategy for dealing with the gap in the case of visual spatial experience is offering a straightforward answer as to why spatial phenomenology has the specific nature that it does, the story being offered here concerning the difference between colour and spatial phenomenology is an account of why we *should not expect an equally satisfying answer* as to why colour phenomenology has the specific nature that it does. In other words, whereas a sense-of-mystery-*removing* explanation is being offered for the specific phenomenal nature of spatial experience, a sense-of-mystery-*preserving* explanation is being offered in the case of colour. A comparison with another sense-of-mystery-preserving theory about consciousness might be helpful here.

The ‘phenomenal concept strategy’¹⁹ seeks to explain why there should be the recalcitrant, illusory *appearance* of an ontological gap between neural/physical properties and phenomenal properties, even though they are in fact identical. This strategy appeals to a distinctive kind of phenomenal concept that we have for phenomenal properties, a concept which can normally only be acquired by actually having the right kind of conscious experience. This phenomenal concept in fact refers to the very same neural process/property that is picked out by a scientific/theoretical concept in neuroscience; but the phenomenal property does so via a very special kind of sensory mode of presentation. The idea then is that due to this special conceptual architecture, we have the unshakeable feeling of there being a genuine explanatory gap when we think of conscious experience first using a phenomenal concept and then using a scientific-theoretical concept, even though these two kinds of concept co-refer (and even if we are committed physicalists!).

The phenomenal concept strategy is thus a sense-of-mystery-preserving theory in that it explains why we should *expect* a persistent (illusory) feeling/appearance that the phenomenal property and the physical property could come apart and are only contingently and arbitrarily related, when in fact they are identical (there is only the one same property being thought of in two very different ways). Likewise, the account offered here, concerning the difference between spatial and colour phenomenology, appeals to the specific features of our visual system and of our perceptual actions to explain why we should *expect* a recalcitrant sense that the specific colour phenomenology of being acquainted with a certain kind of surface could have been quite different, whereas we have no such lingering sense of mystery when it comes to being visually acquainted with a certain kind of shape. In particular, I appealed to the far greater range of sensorimotor skills that we can consciously

¹⁹ Some leading proponents of this strategy include: Loar 1990, Papineau 2002, 2007, Levin 2007, Perry 2001. For a collection of essays largely devoted to phenomenal concepts see: Alter & Walter (eds.), 2007.

employ to independently manipulate and vary both sides of the perceptual interaction in the case of spatial experience compared with colour experience.

Talk of ‘sensorimotor skills’ will naturally bring to mind recent Enactivist theories of perception – and in particular the Sensorimotor theories of Alva Noe and Kevin O’Regan²⁰. And so to finish this section I want to briefly discuss what the relation is between the present account and these Enactivist/Sensorimotor accounts.

If we take the core Enactivist thesis about perception to be that perception somehow constitutively involves temporally extended perceptual actions – e.g. moving one’s head and one’s viewpoint, saccades of the eyes, shifts in attention, etc. – it should be clear that there is no obvious entailment between EWA and Enactivism in either direction. Indeed, Andy Clark (2009), one of the first, main proponents of Enactivism and the ‘Extended Mind’, has explicitly repudiated EWA²¹. Nor does the specific account I have offered here concerning the especial ‘mysteriousness’ of colour phenomenology require endorsing Enactivism. One could accept that our sense of a special Explanatory Gap for colour experience (as compared with shape experience) is due to the very limited range of exploratory perceptual actions we can perform with respect to colours, *without* thereby having to accept that perception *constitutively involves* such perceptual actions. Nevertheless, it seems reasonable to think that an Enactivist account of the physical and psychological processes required for perception would offer a very natural *implementation story* for how an EWA-style relation of conscious acquaintance with external environmental features could be realised and sustained – especially in view of the emphasis I laid, back in section 2, on how EWA theorists should treat a perceptual experience as a form of *interaction*. Enactivism and EWA might then reasonably be thought to offer each other at least some degree of (non-entailing) mutual support.

Noë and O’Regan’s ‘Sensorimotor’ theories, unlike some other Enactivist theories, are meant to provide an explanation of the specific phenomenal character of conscious perception. And in fact Noë explicitly characterized his own Sensorimotor theory as ‘Naïve Realist’²²:

“Naive Realism is not an alternative to the actionism I develop in *Action in Perception*. My view is as naive realist as one can hope to be. Perceptual experience is a genuine encounter with the world; but it is not a bare encounter; it is an encounter made possible by our possession and exercise of understanding.” (Noë, 2008, 703)

For present purposes, the most important commitment of the Sensorimotor approach is that this practical understanding and exercise of sensorimotor know-how are not just meant to play an essential role in enabling perception, they are supposed to be partially *constitutive* of the distinctive phenomenology of a sense-modality, such as vision. For example, Noe writes:

“The enactive view...proposes that perceiving is constituted by the exercise of a range of sensorimotor skills. ... I want to suggest that one experiences [an object] as having a certain perspectival shape, and so as affording possible movements, only insofar as, in encountering it, one is able to draw on one’s appreciation of the sensorimotor patterns mediating (or that might be mediating) your relation to it.” (Noe, 2004, 90)

Whilst EWA theories need not be committed to this idea, the alleged constitutive contribution of sensorimotor know-how to phenomenal character could be accommodated easily as one more aspect of the specific way or manner in which perceptual acquaintance takes place²³.

²⁰ See e.g. O’Regan & Noe (2001), O’Regan, Myin & Noe (2006), Noe (2004), O’Regan (2011).

²¹ See Ward (2012) for a reply to Clark.

²² See also Noe (2005) for another explicit statement of his Naïve Realist allegiance.

Notice also that whilst the explanation I have offered here concerning the difference between spatial and colour phenomenology also appealed to the range of sensorimotor skills and implicit know-how that is available in each case, this does not require endorsing Noe & O'Regan's thesis that sensorimotor know-how is constitutive of phenomenal character. Again: one might think that the different ranges of sensorimotor skills that one can consciously exercise in the case of colour perception and (visual) spatial perception is what explains the recalcitrant sense of unexplained mystery about colour phenomenology *without* accepting that the exercise of these sensorimotor skills is partially constitutive of either colour or spatial phenomenology. Moreover, the present account turns on there being a far *greater range* of sensorimotor skills and implicit practical knowledge that are typically employed in visual spatial perception compared with only a very impoverished range in the case of colour perception. This emphasis on the *differential importance* of sensorimotor skills in spatial experience as opposed to colour experience is something that is absent from Noe's and O'Regan's theories. Nevertheless, I would hope that the account I have offered here might be still be congenial to sensorimotor theorists. For example, Kevin O'Regan (2011), one of the very few psychologists to explicitly address the "why is it specifically like this?" question, offers an explanation for why the specific phenomenal character of experience, and in particular colour experience, can be 'ineffable':

"...the ineffability of feels is a natural consequence of thinking about them in terms of ways of interacting with the environment. Feels are qualities of sensorimotor interactions in which we are currently engaged. We do not have cognitive access to each and every aspect of these interactions." (O'Regan, 2011, 115)

"The sensorimotor approach cannot account for all the aspects of the experience of colour, because sensory experiences are by nature ineffable – their aspects cannot be completely apprehended within our cognitive structures because we do not have cognitive access to every detail of the skills we exercise when we interact with the world." (ibid. 132)

The story I have told here could be understood as *supplementing* what O'Regan claims in these passages, by focusing on how much more easily we can 'apprehend' the nature of our visual interactions with shapes/spatial properties compared with our visual interactions with colour, and so how much less ineffable or mysterious the specific phenomenal character of spatial experience seems to us compared with colour experience.

8. Summary & Conclusion

It is often said that what is distinctively difficult about the Explanatory Gap, what is so hard about the 'Hard Problem of Consciousness', is that we struggle to imagine what form a physicalist/functionalist solution could possibly take²⁴. This is, I take it, a large part of the motivation for the various theorists who claim that the Explanatory Gap is either some kind of cognitive illusion or is in principle insoluble by physicalist/functionalist means. What would a straightforward physicalist solution to the Explanatory Gap even look like?

I have suggested that adherents of EWA who want to address the notorious Explanatory Gap should begin by focusing on shape as the clearest case, and then go on to explain why we should not expect colour to be so easy. If visual spatial perception is understood in terms of EWA, then we have at least one case where the Specific What-it's-Like Question can be intelligibly answered (given some viable physicalist/functionalist answer to the Why-is-it-like-*anything* question). This is, I suggest, a

²³ Or to put this in Campbell's (2007, 2009) terms: sensorimotor know-how could be thought to be one more factor to include as making up the conscious subject's specific 'viewpoint', the third relatum in Campbell's 3-place acquaintance relation.

²⁴ See Chalmers (1995) for a canonical statement of this idea.

very significant theoretical virtue for EWA and a result that merits the extended discussion I have given it here. For recall (as we saw in section 4) the Specific What-its-Like Question was taken, by such prominent theorists as Jaegwon Kim, Michael Tye and Andy Clark, to be the part of the Explanatory Gap that was most stubbornly resistant to physicalist/functionalist solution. There is, according to the story presented here, no great metaphysical difference when it comes to colour – there is no deep primary quality vs. secondary quality difference between shape perception and colour perception. The reason why the same form of explanation seems less satisfying in the case of colour than with shape is largely a difference between the macroscopic vs. microscopic features we are gaining visual awareness of and the consequently different kinds of *conscious perceptual actions* that can be performed. On both sides of the perceptual interaction when we perceive colours – microscopic, light-reflecting surface texture on the one hand and opponent processing by ganglion cells on the other – the crucial elements are *not* familiar macroscopic features and it is either hard or practically impossible for us to vary them independently and at will whilst consciously observing how they combine.

Of course, none of the foregoing tells us anything about cases of hallucination or dreams, nor how the approach might be extended to other sense-modalities, to bodily sensations, conscious thoughts, emotions etc. Nor is it intended to persuade anyone who thinks that arguments from hallucination/illusion are decisive against any kind of EWA. But insofar as it allows for the *only* example we currently have of how the explanatory gap could be given an intelligible solution, even those who have no sympathy for EWA ought to allow that it has at least this one very interesting feature. For even if you don't buy EWA, the fact that it would allow for an intelligible, physicalist solution to the explanatory gap, at least in one case, shows us something about what a straightforward answer to the gap *would look like*²⁵.

References

- Allen, K. (2016) *A Naïve-Realist Theory of Colour*, Oxford University Press.
- Alter, T., & S. Walter (eds.) (2007) *Phenomenal Concepts and Phenomenal Knowledge. New Essays on Consciousness and Physicalism*, Oxford: Oxford University Press.
- Armstrong, D. (1969) "Color-Realism and the Argument from Microscopes", in Brown and Rollins (eds.) *Contemporary Philosophy in Australia*, London: Allen and Unwin.
- Bermudez, J. (2000) 'Naturalized Sense Data', *Philosophy and Phenomenological Research* 61 (2):353 - 374
- Brewer, B. (2011) *Perception and its Objects*, Oxford: OUP.
- Byrne, A. (2001) "Intentionalism Defended", *Philosophical Review* 110 (2): 199-240.
- Byrne, A. (2003) "Consciousness and Nonconceptual Content", *Philosophical Studies* 113 (3): 261-274.
- Byrne, A. (2006) 'Color and the Mind-Body Problem', *Dialectica*, 60 (2):223-44.
- Byrne, A. & Hilbert, D. (2003) "Color Realism and Color Science", *Behavioral and Brain Sciences* 26 (1): 3-21.
- Byrne, A. & Tye, M. (2006) 'Qualia ain't in the head' *Nous* 40(2): 241-255.
- Campbell, J. (2002) *Reference and Consciousness*, Oxford: OUP.
- Campbell, J. (2005) "Transparency vs. Revelation in Color Perception", *Philosophical Topics*, 33: 105–115.
- Campbell, J. (2007) "The metaphysics of perception", *Philosophical Issues* 17: 1-15.

²⁵ Versions of this paper were presented at events at the University of Leeds, King's College London, the University of Barcelona, the Institute of Philosophy in London and the University of Antwerp. Many thanks to the audiences on all those occasions. Thanks in particular to Dominic Alford-Duguid, Marc Artiga, Mike Beaton, Ori Beck, Bill Brewer, Tony Cheng, Ellen Fridland, Anil Gomes, Jonathan Knowles, Heather Logue, Clare Mac Cumhaill, Mike Martin, Anders Nes and Stephen Raleigh for helpful comments and criticisms. Finally, I am very grateful to two anonymous referees for this journal for their extremely perceptive and helpful reports, which greatly improved this paper.

- Chalmers (1995) “Facing up to the problem of consciousness”, *Journal of Consciousness Studies* 2: 200–19.
- Clark, A. (2000) “A case where access implies qualia?” *Analysis* 60: 30-37.
- Clark, A. (2009) ‘Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head’, *Mind* 118: 963-993
- Cohen, J. (2009) *The Red and the Real: An Essay on Color Ontology*. Oxford: OUP.
- Dretske, F. (1996) ‘Phenomenal Externalism or If Meanings Ain’t in the Head, Where Are Qualia?’ *Philosophical Issues* Vol. 7 (Perception): 143-158.
- Fish, W. (2008) “Relationalism and the Problems of Consciousness”, *Teorema* Vol. XXVII/3: 167-180
- Fish, W. (2009) *Perception, Illusion and Hallucination*, Oxford: OUP.
- Gert, J. (2008) ‘What Colors Could Not Be: An Argument for Color Primitivism’, *Journal of Philosophy* 105: 128–57
- Genone, J. (2016) ‘Recent Work on Naïve Realism’, *American Philosophical Quarterly* 53 (1).
- Hacker, P (1987) *Appearance and Reality: A Philosophical Investigation into Perception and Perceptual Qualities*, Oxford: Blackwell.
- Hinton, J. (1967) “Visual Experiences”, *Mind*, 76(April): 217–227.
- Hinton, J. (1973) *Experiences: An Inquiry into Some Ambiguities*, Oxford: Clarendon Press.
- Jackson, F. (1977) *Perception: A Representative Theory*, Cambridge University Press.
- Jackson, F. (1982) “Epiphenomenal Qualia” *Philosophical Quarterly* 32: 127-136.
- Jackson, F. (1998) *From Metaphysics to Ethics: A Defence of Conceptual Analysis*, Oxford: OUP.
- Johnston M. (2006) “Better than mere knowledge? The function of sensory awareness” in *Perceptual Experience*, ed. T. Gendler & J. Hawthorne, Oxford: OUP.
- Kalderon, M. (2007) ‘Colour Pluralism’, *Philosophical Review* 116: 563–601.
- Kennedy, M. (2007) “Visual Awareness of Properties” *Philosophy and Phenomenological Research* 75 (2): 298-325
- Kim, J. (2007) *Physicalism: Or something near enough*. Princeton University Press
- Kirk, R. (2005) *Zombies and Consciousness*. Oxford: OUP.
- Langsam, H. (2011) *The Wonder of Consciousness*. Cambridge: MIT Press
- Levin, J. (2007) “What is a phenomenal concept?” in T. Alter & S. Walter (eds.) *Phenomenal Concepts and Phenomenal Knowledge. New Essays on Consciousness and Physicalism*, Oxford: Oxford University Press.
- Loar, B. (1990) “Phenomenal States” (Revised Version), in *The Nature of Consciousness: Philosophical Debates*, N. Block, O. Flanagan, G. Güzeldere (eds.), Cambridge, MA: MIT Press.
- Logue, H. (2012) “Why Naïve Realism?” *Proceedings of the Aristotelian Society* 112: 211-237.
- Lycan, W. (2001) “The case for phenomenal externalism”, *Philosophical Perspectives* 15:17-35.
- Martin, M. G. F. (2002) “The Transparency of Experience”, *Mind and Language*, 17: 376–425.
- Martin, M. G. F. (2004) “The Limits of Self-Awareness”, *Philosophical Studies*, 120: 37–89.
- Martin, M. G. F. (2006) “On Being Alienated”, in *Perceptual Experience*, Tamar S. Gendler and John Hawthorne (eds.), Oxford: Oxford University Press.
- McDowell, J. (1982) “Criteria, Defeasibility and Knowledge”, *Proceedings of the British Academy*, 68: 455–79.
- McDowell, J. (2008) “The disjunctive conception of experience as material for a transcendental argument”, in *Disjunctivism: Perception, Action, Knowledge*, ed. A. Haddock and F. Macpherson. Oxford: OUP.
- McGinn, C. (1996) ‘Another Look at Color’, *Journal of Philosophy* 93(11): 537-53
- Noe, A. (2004) *Action in Perception*, MIT Press.
- Noe, A. (2005) ‘Real Presence’, *Philosophical Topics* 33 (1): 235-264.
- Noe, A. (2008) ‘Reply to Campbell, Martin and Kelly’, *Philosophy and Phenomenological Research*, Vol. 76, No. 3: 691-706
- O’Regan, K. (2011) *Why Red Doesn’t Sound Like a Bell*, Oxford University Press.
- O’Regan, K., Myin, E. & Noe, A. (2006) “Skill, Corporality and Alerting Capacity in an Account of Sensory Consciousness” in Laureys, S. (ed.) *Boundaries of Consciousness*, Elsevier.

- O'Regan, K. & Noe, A. (2001) 'A sensorimotor approach to vision and visual consciousness' *Behavioural and Brain Sciences*, 24(5): 939-973.
- Palmer, S. (1999) *Vision Science: Photons to Phenomenology*, Cambridge, MA: MIT Press.
- Papineau, D. (2002) *Thinking about Consciousness*, Oxford: Oxford University Press.
- Papineau, D. (2007) "Phenomenal and perceptual concepts," in T. Alter & S. Walter (eds.) *Phenomenal Concepts and Phenomenal Knowledge. New Essays on Consciousness and Physicalism*, Oxford: Oxford University Press.
- Papineau, D. (2014) "Sensory Experience and Representational Properties", *Proceedings of the Aristotelian Society* 114 (1): 1-33.
- Perry, J. (2001) *Knowledge, Possibility and Consciousness*, Cambridge, MA: MIT Press.
- Prinz, J. (2006) "Putting the Brakes on Enactive Perception" *Psyche* 12 (1)
- Putnam, H. (1975) "Philosophy and Our Mental Life", in *Mind, Language, and Reality*, Cambridge University Press.
- Raleigh, T. (2015) "Phenomenology without Representation", *European Journal of Philosophy*, 23(4): 1209-1237.
- Ratliff, F. (1965) *Mach Bands: Quantitative studies on neural networks in the retina*, San Francisco: Holden-Day.
- Reid, T., 1822/1970, *An Inquiry into the Human Mind*, (ed.) T. Duggan, Chicago: Chicago University Press.
- Ross, P. (forthcoming) 'Phenomenal Externalism's Explanatory Power', *Philosophy and Phenomenological Research*
- Russell, B. (1913) "The Nature of Sense-Data –A Reply to Dr Dawes Hicks", *Mind* 22 (85):76-81.
- Russell, B. (1918) *Mysticism and Logic and Other Essays*, New York, London: Longmans, Green & Co.
- Shoemaker, S. (1975) 'Functionalism and Qualia', *Philosophical Studies* 27: 291-315.
- Snowdon, P. (1981) "Perception, Vision and Causation", *Proceedings of the Aristotelian Society*, 81: 175–92.
- Soteriou, M. (2013) *The Mind's Construction*, Oxford: OUP.
- Soteriou, M. (2016) *Disjunctivism*, Routledge.
- Travis, C (2004) "The Silence of the Senses", *Mind*, 113(449): 57–94.
- Tye, M. (2000) *Consciousness, Color, and Content*, Cambridge, MA: MIT Press
- Tye, M. (2006) "Absent Qualia and the Mind-Body Problem" *The Philosophical Review* 115: 139-168.
- Tye, M. (2009) *Consciousness Revisited: Materialism Without Phenomenal Concepts*. Cambridge, MA: MIT Press.
- Tye, M. (2015) "Phenomenal Externalism, Lolita, and the Planet Xenon," in *Qualia and Mental Causation in a Physical World: Themes from the Philosophy of Jaegwon Kim*, eds. Terence E. Horgan, Marcelo Sabates, and David Sosa. Cambridge: Cambridge University Press, 190–208.
- Ward, D. (2012) 'Enjoying the Spread' *Mind* 121 (483): 731-751.
- Wright, C. (2008) 'Comment on John McDowell's "The Disjunctive Conception of Experience as Material for a Transcendental Argument"', in *Disjunctivism: Perception, Action, Knowledge*, ed. A. Haddock and F. Macpherson. Oxford: OUP.