Sensorimotor knowledge and the contents of experience

Julian Kiverstein
School of Philosophy, Psychology and Language Sciences
University of Edinburgh
Dugald Stewart Building
3 Charles Street
Edinburgh
Scotland, EH8 9AD

Email: j.kiverstein@ed.ac.uk

Sensorimotor knowledge and the contents of experience

Introduction

Perhaps the defining feature of perceptual experiences is its intentionality or directedness upon the world. As I look about my office I can see pictures on the wall, my houseplant and the books and papers scattered around my messy desk. My visual experience is of, or about, these things by virtue of its intentional content, but in virtue of what do our experiences have intentionality? When two perceptual experiences have the same intentional content or different intentional content what is it that explains this sameness or difference?

I will argue that at least some aspects of the intentional content of experience depend on lawlike relationships between sensory input and motor output. Hence I will be defending a weak version of the sensorimotor theory of consciousness.[1] According to this theory, perceptual experiences are skilled activities in which bodily movements are used to probe and explore the environment. Perceivers understand how what they perceive systematically depends on what they do. I will refer to this understanding as *sensorimotor knowledge.* What role does sensorimotor knowledge play in determining an experience's intentional content? This is the question I will take up in what follows.

Sensorimotor knowledge consists in a perceiver's familiarity or attunement with the lawlike ways in which sensory stimulation varies with movement.[2] O'Regan and Noë (2001) refer to these laws as *sensorimotor contingencies* (following MacKay (1962)) and describe two varieties of such laws.[3] There are laws that relate to the distinctive kind of sensory apparatus we have, and the ways in which motor actions produce sensory changes in creatures with

[1] See Merleau-Ponty 1962; Hurley 1998; O'Regan & Noë 2001; Noë 2004; Siewert 2005; and Thompson 2007, for different versions of this idea. The sensorimotor theory also goes under the name of the "enactive" theory of perception (Noë 2004; Thompson 2007) and actionism (Noë 2007).

[2] Clark (2006) has pointed out that there is an ambiguity in this way of describing sensorimotor knowledge. Sensory input could be understood in personal-level terms as something like the conscious sensations a perceiver undergoes or it could be understood in sub-personal terms as raw, unprocessed, proximal sensory stimulation. I take up this issue in §1.

[3] In his single-authored work, Noë distinguishes between "movement-dependent" and "object-dependent" sensorimotor laws (2004, p.64). Movement dependent laws describe the systematic ways in which movements of the body modulate sensory stimulation. Object-dependent laws on the other hand describe how movements of the object would produce changes in sensory stimulation. To perceive something, Noë claims a perceiver must draw on her understanding of both kinds of laws.

sensory apparatus like ours. Consider how the sensory stimulation the eye receives changes in very particular ways determined by the spherical shape of the retina. As you fixate on the midpoint of a horizontal line "the line will trace out a great arc on the inside of your eyeball. If you now switch your fixation point upwards, the curvature of the line will change" (O'Regan & Noë 2001: 941). Straight lines distort as we move our eyes in regular ways that are determined by the shape of our eyes. Having a visual apparatus like ours will determine a set of laws that relate systematic changes in sensory stimulation to motor behaviour.

A second class of contingencies relate to how the sensory attributes of an object are presented under different viewing conditions. Philopona & O'Regan (2006) describe a set of laws that determine how coloured surfaces transform incoming light into outgoing light. As one moves a red piece of paper around under different illuminations the light the paper reflects will be different depending on whether the paper is receiving bluish light from the sky, yellowish light from the sun, or reddish light from a streetlight. The sensorimotor contingencies that govern our experience of a surface's colour are laws that determine how the light will reflect off the surface of objects as they are moved around under different illuminations.[4]

Previous sensorimotor theories have made the strong claim that the exercise of sensorimotor knowledge can fully explain the content and character of our experiences. I will label this the "strong sensorimotor hypothesis" (SSH). Hurley & Noë (2003) have argued for SSH by describing how the sensorimotor theory can address what they call the "comparative explanatory gaps". It can explain why cortical neural activity should realise experiences that instantiate one type of quality (e.g. redness) rather than a distinct type of quality (e.g. greenness). It can also explain why neural activity of a certain type should realise experiences in a particular sense modality (e.g. sight rather than touch). The key to their argument is to look at cases in which neural activity changes its "qualitative expression" from realising experiences of one type to realising experiences of a distinct type. If we can explain why a change of this type takes place, Hurley & Noë reason that this will amount to answering the questions the comparative gap raises. They proceed to argue that:

[4] Philopona & O'Regan found that objects colured red, yellow, blue and green all affect incoming light in simpler ways than other colours. From this finding they deduced a number of well-established facts from colour science such as facts about colour naming, unique hues and hue cancellation.

> "…changes in qualitative expression are to be explained not just in terms of the properties of sensory inputs and of the brain region that receives them, but in terms of dynamic patterns of interdependence between sensory stimulation and embodied activity. What drives changes in qualitative expression of a given area of cortex….(are) higher-order changes, in relations between mappings from different sources of input to different areas of cortex and from cortex back out to effects on those sources of input, which are in turn fed back to various areas of cortex." (Hurley & Noë 2003: 146)

To defend SSH it would have to be shown that there are patterns of interdependence between sensory input and motor behaviour that explain both the differences between sight as contrasted with touch, and also the differences between seeing one quality rather than another. Noë has for instance appealed to what he calls "sensorimotor profiles" to account for our experience of shape.[5] An object's sensorimotor profile determines how this object's appearance would undergo transformation as one varies the location from which the object is seen. If it could be shown that every quality of experience can be explained in this way in terms of sensorimotor profiles, we would have an answer to the question with which we started. We could explain when we have experiences with the same content and when we have experiences with different contents by appealing to the sensorimotor profiles that govern the respective experiences.

Clark (2006 & 2008, ch.8) has however argued that SSH is incompatible with the two visual systems hypothesis (Milner and Goodale (1995/2006); Goodale and Milner 2004) according to which humans and primates have two distinct visual systems. Following Clark, I will take the two visual systems hypothesis to claim that a computationally efficient coding of visual information for reasoning, recall and planning, precludes the use of the same coding for the online control of action. The online control of action seems to require the use of a very different kind of visual information from that which is used for reasoning, recall and planning. The use of visual information to guide and control action requires visual representations that are constantly updated, egocentrically specified, and sensitive to distance and orientation. While the visual information used in reasoning, recall and planning is by contrast object-

[5] O'Regan & Noë (2001, 942) made a similar claim. They write: "The idea we wish to suggest here is that the visual quality of shape is *precisely* the set of all potential distortions that the shape undergoes when it is moved relative to us, or when we move relative to it. Although this is an infinite set, the brain can abstract from this set a series of laws, and it is this set of laws that codes shape."

centred (as opposed to egocentric) and items are represented by their category and significance irrespective of differences in retinal input, location, and viewpoint. Vision is clearly involved in the performance of both sets of tasks, but the visual representations involved in the two cases seem to have very different kinds of content. Hence the hypothesis that there are two visual systems one dedicated to vision for perception and the other to vision for action.[6]

Clark has argued that one important moral of the two visual systems work is that the contents of *conscious* vision are "tweaked and optimised" in ways that do not "march in step" with sensory stimulation (2008: 179).[7] The contents of conscious experience are suited to tasks like planning and selecting types of action or identifying and discriminating an object's properties. These tasks do not require representations that are sensitive to every nuance in sensory stimulation and the changes that take place in sensory stimulation with movement. What will prove important for success in these tasks is locking onto features such as rough spatial location, colour, size, shape and so on, information that is useful for selecting targets for action and types of action. This conclusion is, he argues, incompatible with SSH which seeks to explain differences in content in terms of sensorimotor contingencies or the lawlike ways in which sensory stimulation changes with movement.

---

[6] Computational considerations of the kind I have just been appealing to are, of course, just one source of evidence marshalled to support the two visual systems hypothesis. Additional evidence comes from double dissociation of vision for perception and vision for action and from a range of experiments that seem to establish that visuomotor behaviour can be accurate when subject's undergo illusory visual experiences of shape or size. Jacob & Jeannerod (2003, part II) offer a detailed and careful discussion of a wide range of evidence for the two visual systems hypothesis.

[7] Clark offers as an example the ingenius experiments by Aglioti et al (1995) which seem to seem to show that our reaching and grasping actions are immune to the Titchener circles or Ebbinghaus illusion (figure 1). He argues that the illusory experience of disk size in this illusion is the result of the conscious visual system making use of assumptions that preserve size constancy. Aglioti et al showed that the motor system produces an aperture grip size perfectly suited to the size of the disks despite the subject's illusory experience of the disk's size. This seems to show the processing underlying visual awareness operates according to abstract assumptions that do not reflect the fine details of our sensorimotor interactions with the world.

Hurley and Noë (2007, §7) respond that the sensorimotor skills that explain in this case the visual experience of the disk's size go well beyond the skills requires for reaching and grasping in this particular experimental context. The skills in question are intended to explain our experience of the disk's size across different actions and context. It is one thing to show that the particular skills required for reaching and grasping a disk in a particular context aren't affected by a subject's illusory experience. It is quite another to show that a significant part of the broad set of sensorimotor skills that explain our experience of size and shape could persist while the subject nevertheless continues to undergo an illusory experience.

I am going to assume initially that the worry Clark is raising here about SSH is well founded. I will begin by arguing that there are nevertheless two ways in which sensorimotor knowledge can contribute to the contents of experience.[8] Hence even if SSH is challenged by the evidence for two visual systems, still a weaker form of the sensorimotor theory may survive. In the final part of my paper I tentatively suggest that once we see the two ways in which sensorimotor knowledge contributes to the contents of experience a defence of SSH may also open up.

§'s 1 and 2 offer some preliminary clarifications of how I think we should understand sensorimotor knowledge. The main argument of the paper begins in §3, where I show how sensorimotor knowledge can provide us with an experience of what forms the periphery or background to our experience. I will defend the view that the contents of experience go beyond what we are attending to at a given time to include items that are unattended.[9] I will argue that sensorimotor knowledge accounts for the perceiver's grip on those unattended items.

In §4 I argue that sensorimotor knowledge can make a second contribution to perceptual content by delivering what I will call "experiential constancy". We experience an object's properties as staying the same as we vary our relation to an object, or as the object is presented under different viewing conditions. We do so by forming what I will call sensorimotor expectations: expectations about how the object would show up under different viewing conditions, or as we vary our relation to the object. When these expectations are correct we are able to keep track of a property across variations in experience. However, we form these sensorimotor expectations by exercising our sensorimotor knowledge. Thus it is sensorimotor knowledge that explains how it is that our experiences can exhibit constancy.

I then return to Clark's worry about SSH. Recall that Clark's objection to SSH was that conscious visual representations are tweaked and optimised so that they can be used, amongst other things, for the selection of targets for action. If visual experiences are to serve this function they must have contents that

[8] In what follows I will concentrate almost entirely on visual perception, though I think most of the claims I make could be extended to the other senses.

[9] Schwitzgebel (2007) draws a contrast between views that claim the contents of consciousness are "rich" and can include unattended objects and views that claim that the contents of consciousness are "thin" and only include what we are currently attending to. I will defend a rich view.

exhibit constancy. Clark argues that the visual system produces such experiences by abstracting away from the fine details of our sensorimotor engagement with the world. However in §4 I will argue that constancy is something we experience by virtue of our sensorimotor expectations. If I am right it is only by employing our familiarity with laws of sensorimotor contingency that we can experience constancy. Hence the representations the visual system uses in reasoning, recognising and planning don't abstract away from our sensorimotor engagement with the world. Rather they depend on our sensorimotor knowledge.

1. Sensorimotor contingencies: proximal or distal?

I have defined "sensorimotor knowledge" as implicit or practical knowledge of sensorimotor contingencies, the laws that determine how "sensory stimulation" varies with motor behaviour. Clark (2006) has argued that the notion of sensory stimulation appealed to in this definition can be understood in at least ways. According to the first interpretation, sensory stimulation refers to raw, unprocessed sensory input such as light received at the retina. While the second interpretation claims that sensory stimulation refers to conscious sensations – states that there is something it is like for a perceiver to undergo. Which, if either, of these options presents the correct interpretation of sensory stimulation?

Briscoe (2008) has raised a related question. He asks whether sensory stimulation should be understood proximally or distally. Understood in the first way we get the two options we have seen Clark describe. Understood in the second way, sensorimotor contingencies are laws that describe changes in a perceiver's sensory relation to the world that are brought about through movement.

A third question we must settle asks whether we should conceive of sensorimotor knowledge as a notion to be deployed in personal level or sub-personal level explanation. Very roughly, the distinction between personal-level and sub-personal level explanation can be mapped onto the distinction between contents and vehicles. Sub-personal explanations describe functional or neural mechanisms that carry content and can be deployed in causal explanations of behaviour. Personal-level explanations don't deal with vehicles but with the contents of a mental state. A mental state will by virtue of its content enter into normative and rational relations with other mental states.

These normative and rational relations will determine how the mental state figures in rational explanations of behaviour. The third issue we must settle asks whether sensorimotor knowledge is to be understood as a notion deployed in explaining vehicles or contents.

In what remains of this section I will attempt to decide the issue of whether sensory stimuli should be understood proximally or distally. If it should turn out that sensory stimulation must be understood distally, this will immediately rule out Clark's two possibilities both of which take sensory stimuli to be something proximal. In §2 I take up the question of whether sensorimotor knowledge is a personal-level or sub-personal level concept.

Sensorimotor theories share in common with ecological theories of perception the claim that we can *directly* perceive external objects and their properties. (By an "external object" I mean a spatially located entity that could in principle exist independent of being perceived.) I shall argue that for this reason we shouldn't understand sensorimotor contingencies in proximal terms. To do so would be to introduce an intermediary into perception, and depending on how we construe this intermediary (as a conscious sensation or as raw, unprocessed sensory input) the result is a theory of perception that is incompatible with the claim that we can directly perceive external objects. Before I turn to this argument we first need to know what is meant by direct perception. Second it must be shown that sensorimotor theories are committed to the claim that direct perception is possible.

I'm going to adopt Paul Snowdon's (1992) definition of direct perception (d-perception):

> "x d-perceives y iff x stands, in virtue of x's perceptual experience, in such a relation to y that, if x could make demonstrative judgements, then it would be possible for x to make the true demonstrative judgement "that is y"." (Snowdon, 1992: 56)

Suppose I am currently looking at a glass of water on my desk. If I were to now make the demonstrative judgement *that* glass is full based on my experience, I could be singling out an external object – the glass before me – or I could be referring to a mind-dependent object, an "image" or "idea" of the glass. To claim that we can directly perceive external objects is to claim that in making a

demonstrative judgement I single out an external object, and not a mind-dependent object. Snowdon's reading of directness doesn't equate direct perception with demonstrative judgement. Rather he uses demonstrative contact with an external object to clarify what is meant by "directness". If perception weren't direct, our demonstrative judgements would never achieve contact with external objects; this contact would always proceed via some mind-dependent intermediary. It is the directness of our perceptual experience that explains our ability to demonstratively think about particular external objects (c.f. Campbell 2002).

I have claimed that sensorimotor theories are committed to the claim that we can directly perceive external objects. The sensorimotor theory claims that perception is a skilled activity in which we use our bodies to probe and explore the external objects around us. The world is used as its own best model (Brooks 1991) and is sampled and interrogated on a need to know basis. The visual system might well produce global representations that give the perceiver the gist of a scene thereby assisting the further exploration of useful locations. It doesn't however go to the effort of producing detailed internal model of the external world. Instead the visual system makes use of bodily movements to instantaneously access the location in the world where the required information is to be found. I suggest that this process of accessing information in the world is precisely the direct perception of external object, which we have seen makes possible true demonstrative judgements.

It might naturally be objected at this point that we only get to perceive the lay of the land around us by transforming an impoverished retinal image into a representation of 3-d scene. Didn't Marr teach us that the task facing the visual system is one of "inverse optics" inferring from impoverished retinal image an internal model in which the properties of visual scene are represented? Consider a real world visual scene such as a densely cluttered forest. The scene will contain all manner of changes in illumination that do not correspond to the contours of the trees and leaves visible in the scene. How do we get to perceive the leaves from the trees if not by means of some fancy inferences carried out by the visual system?

Sensorimotor theories deny that any such inference is necessary. In common with ecological theories of perception, they claim that movement can resolve many of the ambiguities in visual input. A large tree that is far away might

produce the same retinal input as a small tree close up. The sequence of sensory inputs generated by movement can resolve the question of whether the tree the perceiver sees is large or small. Many of the questions that motivate the inverse optics research program assume that the visual system operates with a stable and static retinal image. If we accept this assumption, it follows that any ambiguities must be resolved by downstream processing. The activities of looking and moving are conceived of as a means of sampling information that only later gets transformed into a visual representation. However we now know from research on active vision that the assumption that the visual processing begins with a stable and static retinal image is almost certainly false.[10] Looking and moving are fundamental parts of visual processing.

It is because the visual system doesn't face a poverty of stimulus problem that perception can be direct. The visual system doesn't start out with ambiguous, impoverished sensory stimuli on the basis of which it faces the task of constructing an internal model of the perceiver's environment. Sensorimotor theories, in common with ecological theories of perception, claim that movement can reveal higher-order relationships between sensory inputs that "specify" to a perceiver what is where. It is by locking onto these higher-order or lawlike relationships in the flow of sensory input that the perceiver is able to extract information about how things are in her surrounding environment. Sensorimotor theories go further than ecological theories in holding that there are also higher-order or lawlike relationships that hold between sensory input and motor output.[11] It is the perceiver's *sensorimotor knowledge* – her familiarity with or understanding of these laws – that allows her to escape any poverty of stimulus problem.

Let us return now to the question of whether sensorimotor contingencies are laws that govern changes in *proximal* sensory stimulation to movement. We should be able to see now why neither of the interpretations Clark offers can be satisfactory. Suppose we take sensory stimuli to refer to conscious sensations that vary with movement. If this were the correct interpretation, whenever we make a demonstrative judgement based on our experience we would be referring not to an external object. In saying "*that* glass is full" based on my visual experience of the glass, I would be referring to some conscious

[10] For an excellent, comprehensive review see Findlay & Gilchrist (2003).

[11] Hurley 2001 explores this difference in more detail than I can go into here. For further discussion also see Gangopadhay & Kiverstein 2009.

sensations that may, or may not, have been caused by the glass. My conscious experience would not place me in any kind of relation to an external, space occupying glass. This consequence conflicts with the sensorimotor theory's commitment to the directness of perception. Hence it cannot be right to claim that sensorimotor contingencies are laws that relate changes in conscious sensation to movement.

Could the notion of sensory stimuli appealed to be sensorimotor theory refer to raw, unprocessed sensory input? All theories of perception must accept that perceptual processing begins with raw, unprocessed sensory input impacting the sense organs. We have seen however that there is some disagreement amongst psychologists about the nature of this sensory input. Proponents of the inverse optics research program hold that sensory input radically underdetermines the perceived scene so that some heavy-duty computational machinery has to be brought to bear to arrive at a representation of a scene. Sensorimotor theories by contrast share with ecological theories the assumption that sensory input is replete with information. J.J. Gibson (1972) distinguished between "optical stimulation" and "optical information". Optical stimulation occurs when the photoreceptors in the eye receive light but this light carries no information about the layout of the environment. Optical stimulation is what the eye receives when the perceiver is immersed in a dense fog, or when the perceiver wears plastic diffusing eye-caps that cause him to see nothing but a homogenous field of whiteness, a *Ganzfeld.* Optical information by contrast is what the perceiver normally receives as sensory input. It is light that has bounced off the surfaces in a perceiver's surrounding environment and converged on a point of observation, which Gibson called an "ambient optic array". The light converging on a point of observation stands in a uniquely specifying, lawlike relationship with the surfaces in the perceiver's field of view. Optical information is information in the light that stands in this lawlike relation of "specifying" the substances and surfaces around the perceiver. According to Gibson, and those that have followed him including, I claim, sensorimotor theories, the eye doesn't just receive optical stimulation; it receives *optical information*. Sensory input as it figures in sensorimotor theories must be understood in relation to the substances and surfaces in the environment it specifies. In other words, sensory stimuli cannot be understood solely in proximal terms. I conclude then that sensorimotor

theories don't interpret sensory stimuli in either of the ways Clark proposes.[12] We should understand sensory stimuli distally. This still leaves open the issue of whether sensorimotor knowledge should be understood as a notion to be deployed in explaining contents or vehicles. It is to this issue that I turn in the next section.

2. Sensorimotor Knowledge and the Personal/Sub-Personal Distinction

Alva Noë (2002; 2004; 2007; forthcoming) has argued that seeing is two-step process: we learn how things are from seeing how they look. When we see a round plate from an angle the plate may look elliptical because of our perspective on the plate. We don't suppose for a moment that the plate we are seeing has an elliptical shape. We see that the plate is round but, according to Noë, we also see a plate that looks elliptical. The plate's looking elliptical is what Noë calls a "perspectival property" of the plate. Perspectival properties are features of a mind-independent, objective world, not of our experiences or sensations. They are aspects of things that are visible to a perceiver from a particular vantage point, the location a perceiver's body occupies at the time. The perspectival properties an object has will vary depending on illumination, the location from which the object is viewed, and other facts about the perceiver. As the perceiver changes his perspective, so the perspectival properties the object instantiates will also vary. As we move we explore how things appear and in doing so we discover how they are:

> "To discover how things are, from how they appear, is to discover an order or pattern in their appearances. The process of perceiving, of finding out how things are, is a process of meeting the world; it is an activity of skillful exploration." (Noë 2004: 164)

Noë describes sensorimotor knowledge as a perceiver's understanding of how an object's perspectival properties normally vary with movement. We have, for instance, a grasp of the way the aspect or profile a plate presents changes when we move. Noë says we encounter "the actual shape of the plate in thus

[12] Clark offers his two readings of the sensorimotor theory in the context of challenging the sensorimotor theory's claim to provide an adequate explanation of consciousness. He mounts a dilemma for proponents of the sensorimotor theory. If, on the one hand, they understand sensorimotor contingencies as laws relating conscious sensations with movement, the account of experience they give presupposes what needs to be explained. If, on the other hand, sensorimotor contingencies are understood as laws relating raw sensory stimuli to movement, it is hard to see how appealing to laws of this kind can help explain consciousness. I don't mean to claim to have answered this objection, only to challenge Clark's way of setting up the problem. My thanks to Dave Ward for pressing me on this point.

bringing to bear our sensorimotor understanding" (2004: 78). Noë has compared this account of perception with Richard Wollheim's (1968/1980: 205-26) account of picture perception as seeing-in. Wollheim describes how we see-in the arrangement of a paint on the surface of painting the scene that is depicted in the painting. Perspectival properties are supposed to be analogous with a picture-surface. Just as we see-in the picture surface the scene depicted, so we can see-in a perspectival property an object's intrinsic properties. Crucially what allows us to go beyond an object's perspectival properties to see its intrinsic properties is our sensorimotor knowledge. We understand, for instance, how the elliptical perspectival property the plate presents when seen from a particular place changes as a result of movement. It is this understanding that makes it possible for us to see roundness in the plate's elliptical profile.

Noë takes sensorimotor knowledge to be a set of *conceptual skills.* He writes: "it is because of our possession" of these skills "that we can have experiences with world-presenting content" (2004: 183). Without sensorimotor knowledge we would have no perceptual access to the world. The absence of sensorimotor knowledge wouldn't deprive us of perceptual access to the world in the way that a lesion to visual cortex might render a subject unable to see colour or motion. What the perceiver would lack is a means of finding out how things are based on how they appear. The way things appear to a perceiver in experience puts him in a position to judge how things are. This is something experiences can do because of the perceiver's sensorimotor knowledge. If a perceiver lacked sensorimotor knowledge he would have no grip on how things might be. It is sensorimotor knowledge that confers this grip on a perceiver.

Does this make sensorimotor knowledge a personal-level concept? Noë takes sensorimotor knowledge to explain how it is possible for experience to play a rational role and so to figure in personal level explanations of our behaviour. Sensorimotor skills are preconditions for having an experience that can form the basis for judgements. Thus it would seem that for Noë sensorimotor knowledge is what enables perception to play a role in personal-level explanation. It doesn't itself figure in personal-level explanation. Noë however has more to say on this question. He tells us that sensorimotor skills:

> "…may be partially or wholly subpersonal. The patterns of stimulation if nerve fibers on the retina belong not to the perceiver's psychology but rather to the conditions that causally enable that psychology. Moreover, these patterns may occur below the threshold of consciousness." (Noë, 2004: 201)

Earlier in the book he goes as far to suggest that when it comes to a theory of perception there is no "sharp line" between the personal and the sub-personal. This conclusion seems to be required by his claim that sensorimotor skills are on the one hand primitive conceptual skills that make it possible for experiences to have contents that can be taken up in making judgements about the world. Conceptual skills are normally thought of as subject to the normative constraints of rationality and the holism of the mental that are the markers of personal-level explanation. On the other hand, as we have just seen Noë also claims that many sensorimotor skills are both conceptual and subpersonal. Thus he concludes the personal/sub-personal distinction is not a useful one for a theory of perception.

I want to suggest a slightly different way of thinking about sensorimotor knowledge that Noë may or may not want to assent to. On this view sensorimotor knowledge is what allows a perceiver to form what I will call "sensorimotor expectations". The latter are expectations about how the perspectival properties an object instantiates will vary with movement and other changes in viewing conditions. Sensorimotor expectations exist at the personal-level: they are expectations relating to what Noë describes as an object's sensorimotor profile. Sensorimotor knowledge on the other hand is sub-personal notion. It is what makes it possible for a perceiver to form accurate or veridical sensorimotor expectations. More precisely, it is the perceiving organism's familiarity with or attunement to sensorimotor contingencies. I've argued that sensorimotor contingencies are laws relating changes in optical information (as discussed in §1) to movement. Perceptual systems learn to associate movement with changes in optical information. Through this kind of associative learning the perceiving animal acquires sensorimotor knowledge. This knowledge is deployed in the formation of sensorimotor expectations. Armed with this understanding of sensorimotor knowledge I now want to suggest two ways in which it can contribute to the contents of experience.

3. Experience, attention and the background

Our eyes are constantly in motion as we visually experience the world with gaze redirection taking place three to four times per second (Findlay & Gilchrist 2003: 4). Indeed there is compelling evidence that if our eyes were not in continuous motion, blindness could ensue (Ditchburn & Ginsborg 1952; Riggs et al 1953; Krauskopf 1963; Yarbus 1967 (all cited by Noë 2004: 13)). Images stabilised on the retina fade from view in much the same way as the clothes that are continuously on our back or the glasses placed on our nose cease to be felt.

This continuous movement of the eye is no coincidence. As Findlay & Gilchrist (2003: 5) have pointed out bodily movements enable the visual system to combine a restricted (foveal) area of high-resolution processing with an ability to monitor the entire visual field. The retinal image is not homogenous; the retina is, for instance, nearly colour blind outside the high-resolution foveal region. Moreover, projections to the retina away from the central foveal area are given decreased weighting in the processing of visual information. This has the consequence that visual acuity decreases for every degree we move out into peripheral vision. Certain discriminations become impossible for instance when stimuli are presented outside of the central region of the visual field (Findlay & Gilchrist 2003, §2.2.3). In order to demonstrate this progressive decline of acuity Dennett (1991: 53) invites us to try to tell the colour of a playing card held at arms length in the periphery of the visual field. He points out that we won't be able to tell if the card is black, red or a face card. Yet we have the impression of seeing uniform detail right across the visual field from the centre out to the periphery. We don't experience the periphery of our visual field as out of focus and monochrome. We have the impression of seeing a uniformly coloured and detailed visual field because all of the information in our external environment is available to us to be accessed by movements of the eyes, head and whole body. We don't take in all of the information in a visual scene in a single take. We learn about the details in a visual scene through a temporally extended process of looking.

Despite the nonhomogeneity of visual processing it nevertheless seems to us as if we are aware of the densely detailed, stable and persisting environment around us. In order for us to have an awareness of this kind must our experiences actually represent all of this detail?

Not necessarily. To see why not consider what Nigel Thomas (1999) has referred to as the *Refrigerator Illusion.* A child might look in a refrigerator a number of times and conclude that the refrigerator light is on continuously because whenever he has looked the light has been on. The child would of course be mistaken in drawing such a conclusion. Another possibility is that the light is only potentially on when we are not looking. Similarly, it could be that when we are aware of this densely detailed environment around us much of this detail is likewise not actually experienced by us. Just as we find out whether the refrigerator light it is on by looking, so also when we have a question about our surrounding environment we find out the answer by looking.[13] The visual system makes the answer immediately available to us by orienting our gaze to whatever part of a scene we need to interrogate in order to answer our query. The visual system needn't have already encoded the answer before we looked. Much of the dense detail in the environment may figure in our ongoing experience only potentially as something that is available to us and which we know-how to access.

Consider in this light some experiments by McConkie and colleagues (McConkie & Rayner 1975, McConkie & Zola 1979) in which subjects were asked to read from a screen of text, only a small area of which was normal (a few letters to the left of and fifteen to the right of the subject's fixation point). The remainder of the letters on display were continuously perturbed. Subjects nevertheless reported experiencing the whole page as containing normal text. They have an impression of seeing a page of normal text because this is what they are accessing from moment to moment. They are in perceptual contact with a small but moving area of normal text, and this suffices to give them the impression that the page as a whole contains normal text.

How much of the phenomenology of our experience is explicable in this way in terms of potential access to the things and their qualities around us? A surprising amount of what we experience can be explained in this way. Do we experience what is in the periphery of our visual field? We certainly don't experience the periphery in quite the same way as we experience what is in the centre. We do however have some awareness of what lies in the periphery.

[13] O'Regan (1992) makes a similar point in arguing for the "world as outside memory hypothesis" (WOM). WOM claims that visual processing takes place in the world rather than inside the head. Our impression of seeing everything derives from the fact that "if we so much as faintly ask ourselves some question about the environment, an answer is immediately provided by the sensory information on the retina." (p.484)

Consider again my messy desk. Much of the clutter is outside of the central area of my visual field, but what I see isn't confined to what is currently occupying the centre of my visual field. I see mess extending around me. We can also see the occluded parts of an object. The Kanizsa triangle illustrates the phenomena well, (see figure 1) although it is an image and not an object. We experience the hidden parts of one triangle continuing behind the other triangle. Finally, consider our visual experiences as of whole three-dimensional solid, opaque objects such as apple, tables, and books. We always see such objects from a particular perspective or point of view, and not all the parts of such objects are visible to a perceiver from the point of view they occupy at a given time. Yet we undergo experiences that present to us three-dimensional, solid opaque objects. Not everything we see is visible to us.

Figure 1 (insert here)

All of these aspects of our experience constitute what I shall henceforth refer to as an experience's *background*. Everything in the background forms a part of experience by being potentially available to be accessed. However not everything we experience has this character. In order for our experiences to have a background they must also have a *foreground*. The words as they appear on the screen of my laptop currently make up the foreground of my experience. The objects that occupy the foreground of a visual experience change from moment to moment with each shift of our gaze. An experience's foreground will contain whatever objects and properties figures in a perceiver's attentional field at that time. Being a part of a perceiver's attentional field is necessary (but arguably not sufficient) for forming a part of the foreground of experience.[14]

Noë (2004 §2.5 & p.193, 216-7; 2008; forthcoming) has argued that there is no aspect of the phenomenology of our experience that isn't explained by potential access. We might think that the facing sides of objects, such as the screen and keyboard of my laptop, are simply present to us in experience.

[14] I say "close relationship" because not everything we are attending is in the foreground of our experience. I can't currently see the backside of the screen of my laptop. Yet my laptop, backside included, forms a part of what I am currently attending to. Being a part of a perceiver's attentional field doesn't suffice for being a part of the foreground of experience. However being a part of the foreground of experience does suffice for being an object of attention. There is nothing in the foreground of experience that doesn't also form a part of the attentional field. I will have more to say on this point later in this section in my discussion of Noë's claim that perceptual phenomenology is virtual all the way down.

They at least don't show up for us as potentially available to be seen, but are immediately given to us in experience. But, says Noë, "crucially, in an important sense, *nothing* is given. The most that we can have is skill-based access to what is there" (2008: 697). The very same phenomenology attaches to our experience of a particular's facing sides as holds of our experience of the particular as a whole. There is no way, claims Noë, of our taking in all of the details of an object's facing side any more than our experience can encompass all aspects of the object we see. As I look at the label on a bottle I can pay attention to some properties of the label only by ignoring others. There will always be a subset of the visible aspects of an object that do not currently figure in my attention. Noë accepts that there are differences between what is currently visible to me from the point of view I occupy and what is hidden or out of view. However, he claims that the difference between that which is open to view and that which is not, is simply a matter of the ease with which we can access that which is open to view. In perceptually exploring an object we have to do less to access those parts of an object that are visible. We have to work harder to gain access to the hidden aspects of an object.

This strikes me as being a valid description of the foreground, and demonstrates one of the complexities of the relationship between the perceptual foreground and attention. Every object and quality that occupies the foreground also has a background, which is potentially available to us to explore. However Noë's framing of his argument for the claim that "nothing is given" presupposes that there is a phenomenological difference between the foreground and background of an experience. It presupposes that there is a phenomenological difference between what is currently a part of the attentional field and what is currently unattended, but is nevertheless experienced as available to be accessed.[15] It is possible to concede to Noë that

[15] It might be thought that we also need a distinction between the background to experience and what falls outside of the background. (My thanks to Erik Rietveld for pressing this question.) However in order to draw such a distinction we must know how to identify the boundaries of the visual field. It is not obvious to me that the visual field has a sharp boundary as opposed to it being indeterminate how far around our bodies the visual field extends. This was an important theme in Merleau-Ponty (1962) and his discussion of indeterminacy (also see Kelly 2004). A second difficulty is that when we pose the question of where the boundaries of the visual field lie we are presupposing a snapshot conception of experience. We are considering experience at a durationless instant in time and supposing that there will be some determinate answer to the question how far the boundaries of the visual field extend at this time. Both the assumption that the visual field has determinate boundaries and the assumption that there are facts about the character of our experience at a durationless instant in time seem to me mistaken. (I don't mean to attribute either assumption to Rietveld.)

I suspect that what this question is really probing is how to draw the line between what, in a perceptual scene, we experience at a given moment in time and what we do not experience. One can grant that there is a difference between what is experienced but unattended and what

there is always more aspects to the qualities and objects that make up the foreground to experience without giving up on the distinction between foreground and background altogether. In any case as I've already pointed out, Noë's argument for such a phenomenological claim presupposes a distinction between foreground and background.

I propose that we conceive of the foreground to experience as what we are overtly attending to. Fixating an item so that it can be processed by the fovea brings items into the foreground of experience. As perceivers move their eyes, looking first here and then there, so they also move their window of attention about a scene. With each shift in attention the contents of the foreground changes. The background, on the other hand, comprises everything we experience as potentially available to be accessed. In what sense do we have the latter kind of experience?

I suggest (following Noë 2004, §2.5) that the latter kind of experience takes the form of sensorimotor expectations formed on the basis of sensorimotor knowledge. I have expectations about the parts of a scene that form the background to my experience. I anticipate that if I move this will vary my relation to an object bringing into view parts of the object that are currently hidden. Similarly for the objects and qualities that make up the periphery of the visual field. I have a sense of these objects and qualities being available for me to access insofar as I know how to move my body to bring what is in the periphery into view. Similarly for the parts of the objects and qualities that are not open to view either because they are occluded or out of view. I expect my movements, or the movements of the object, or other changes in current viewing conditions to modulate my current experience of an object or quality in particular ways. I experience the hidden parts of an object because my encounter with the object leads me to form a set of sensorimotor expectations. The background is a core aspect of every experience. It forms a part of our experience by virtue of our sensorimotor expectations.[16]

---

is not experienced at all without conceding either of the assumptions described above. How should we mark this distinction? Here I am going to help myself to an answer Noë returns to this question. He claims that objects or qualities are present in experience if our relation to them is mediated by patterns of sensorimotor dependence. There are two kinds of pattern of sensorimotor dependence (as described above in fn.?) and both kinds of pattern have to hold in order for an object or quality to figure in the background of one's experience.

[16] Block (2005) has made a similar suggestion in his review of Noë (2004). (He attributes a related proposal to Sean Kelly.) Block suggests that the background "may be a matter of multi-modal or amodal spatial imagery, and that that imagery may in part be motor imagery – since its brain basis appears to overlap with motor guidance systems in the dorsal visual system" (p.271). What I am calling "sensorimotor expectations" can be thought of in terms of

So far I have argued that the background to experience is best understood as contributing to the contents of experience because we take ourselves to have access to what is currently out of view. In the remainder of the paper I will argue for a second kind of contribution that sensorimotor knowledge can make to the contents of perception. I will show how sensorimotor knowledge can give us *epistemic access* to perceptual constancies. Like the background to our experience perceptual constancies can't be explained in terms of what we are experiencing at a durationless instant in time. The eye doesn't function like a camera giving us access to an object's size, shape or colour by producing a snapshot representation of an object at a discrete instance in time. Our experience of a scene unfolds over time, and we learn about the objects and properties that make up a scene through the exploratory movements of our sense organs. Just like the background to an experience, the objective, intrinsic properties of the things that make up a scene are something we access through exercising our sensorimotor knowledge.

4. Explaining perceptual constancy

Consider again Peacocke's (1983) well-known example of size constancy: you are looking at two trees of the same size, one of which is twice as near as the other. The two trees are represented as being the same size yet the closer tree looks larger than the one further away – the closer tree occupies a larger region in the visual field. On the basis of examples like this one, Peacocke goes on to argue that we must distinguish between representational properties and what he calls "sensational properties". Our experiences have representational properties in virtue of which they represent trees of the same size. The area in a two-dimensional visual field that the trees respectively occupy is a sensational property. Sensational properties should, Peacocke claims, be understood as non-representational properties of experience. If we don't distinguish representational from sensational properties, we wind up with experiences that have contradictory contents. In the current example, we would undergo an experience that represents trees of the same size and also represents trees of different sizes.[17] I will bracket the question of whether Peacocke was right to distinguish sensational properties from representational

amodal spatial imagery. We can think of the formation of sensorimotor expectations in terms of the perceiver simulating movements and their sensory consequences. The latter imagery secures for us an experience of the object's hidden parts.

[17] This consequence can be avoided either by distinguishing different layers of representational content (see Lycan 1995; Tye 2000: 77-9) or by distinguishing situational from intrinsic aspects of the objects of perception (see Schellenberg 2008).

properties. I want to consider instead how the visual system might come to represent the intrinsic, objective size of the trees as contrasted with the size as represented from a particular place?[18]

One way the visual system might solve this problem would be to use information about a perceiver's distance from the trees together with information about the size of the image the trees project onto our retina (arrived at by determining the visual angle the object subtends in relation to the perceiver's eye) to calculate the tree's actual size. By taking into account information about the perceiver's distance from the two trees the visual system could compute the object's size from retinal size.[19] Mistakes occur when the perceiver's distance from an object is miscalculated. If, for instance, the visual system takes the perceiver to be closer to an object than she actually is, the result will be representation of the object as being smaller than it actually is. If the visual system takes the perceiver to be further away from an object than she in fact is the result will be representation of an object as being larger than it actually is.

The same sort of explanation could be given of how our visual experiences come to exhibit other kinds of constancy. As we vary our relation to an object so the proximal stimulation the eye receives will also vary. How does the visual system come to represent the objects invariant or intrinsic shape across these changes in stimulation? Again the visual system could use the retinal image together with information about orientation to compute shape constancy. In a similar manner position constancy could be calculated partly by taking as inputs information about the displacement of the object's image on the retina and information about displacement of the eye relative to the body. Colour constancy could be calculated as a function of the surrounding illumination and the amount of light an object is reflecting.

These computational explanations of constancy explain how the visual system could deliver information about an object's intrinsic properties despite

[18] I take a property *p* to be an "intrinsic" property of an object if *p*'s identity doesn't depend on its bearer's relation to other objects distinct from itself. For further discussion of the distinction between "intrinsic" and "relational" properties see Weatherson 2006.

[19] This of course oversimplifies the sort of explanation one finds of size constancy in vision science. The consensus seems to be that the visual system avails itself of a wide range of information such as information about the relative size of other objects, textural cues such as the number of textural elements occluded, and the relation of an object on a plane to the horizon of that plane. For a discussion of the different factors that might be used to solve size constancy see Palmer 1999, ch.7. §1.1.

variations in proximal stimulation. However this is not quite the problem we were seeking to solve. Consider again Peacocke's example of the two trees: we don't simply see two trees of the same size even though this maybe what our visual systems have computed. The tree that is nearer to us *looks* larger than the tree that is further away: it occupies a larger amount of my visual field than the one that is further away. Consider as a second example, the experience I am currently having of the colour of my office walls. My experience represents the walls of my office as being uniformly coloured white, yet because of variations in the way the sunlight and artificial lighting illuminates the room, and the chromatic properties of surrounding objects some parts of the wall appear darker and other parts appear lighter. These differences in illumination partly determine the way the colour of different parts of the wall appears to me. Suppose I was asked to take part in a psychophysics experiment in which I am asked to adjust the hue and saturation of a colour sample presented on a screen until it matched the way different parts of the wall looked to me. I would make different adjustments in hue and saturation for different parts of the wall. While we are able to recognise the wall's uniform whiteness we can also distinguish the way in which different parts of the wall appear to us. Cohen (2008) has recently described how subjects can even be made to switch between these two aspects of experience by experimental instructions:

> "…instructions to 'adjust the test patch to match its hue and saturation to those of the standard patch' lead subjects to distinguish members of such pairs, while instructions to 'adjust the test patch to look as if it were "cut from the same piece of paper" as the standard, i.e. to match its surface colour' (Arend & Reeves 1986: 1744) lead subjects to assimilate members of such pairs." (Cohen 2008: 66)

The computational accounts of perceptual constancy just described can account for the aspect of our experiences that represents an object's intrinsic properties. Yet we have just seen how this in no way exhausts the phenomenology of our experience. Our experiences represent intrinsic features of an object such as its size, shape, position, and colour as staying the same across variations in appearance. We can simultaneously experiences two instances of a property as the same property whilst also experiencing them as differing in appearance. Similarly we can see an object as retaining identical

features across successive experiences despite the very same features also appearing differently. Certainly my office walls do not appear to change colour as I move my eyes across them, but some regions appear darker, and other regions appear lighter. Any satisfactory account of constancy must do justice to how even given these variations in appearance we nevertheless enjoy experiences that exhibit constancy. The trouble with the computational accounts I sketched above is that they only explain how the visual system might represent information about an object's intrinsic properties. The computational explanation is silent about the aspect of our experience that varies.

Tye (2000, ch.4) has argued that we should think of there being many layers to the contents of experience. Particularly relevant to our current concerns is a distinction he makes between *viewpoint-relative* and *viewpoint-independent* content. Viewpoint-independent content carries information about an object's intrinsic properties, while viewpoint-relative content carries information about an object and its properties relative to some coordinate system centred on the perceiver's body. Consider Peacocke's trees example again: Tye claims that our experience represents the nearer tree as differing in size from the further tree relative to the location of the perceiver's body. The tree that is closer subtends a larger visual angle relative to the eyes of the viewer than the tree that is further away. However there is also a layer of (viewpoint-independent) content that represents the trees as the same size. Tye's account therefore recognises both dimensions of perceptual content we described above. He claims that in every experience a perceiver is aware both of an object's properties relative to a viewing point and of an object's viewer-independent properties.

However, Tye offers no account of how these different layers of content are related to one another. As we move an object around under different illumination the object's colour will appear differently but we don't experience the object as changing its colour. We experience its colour as remaining the same. Tye's suggestion seems to be that we should think of the aspect of the experience that changes as being a distinct layer of experience from the aspect that stays the same. However this is unsatisfactory as can be seen by returning to the sensorimotor theory. As we move our attention about a perceptual scene so the foreground of our experience changes. What is changing here is what Tye labels the "viewpoint-relative" aspect of the content

of experience. Yet across these changes in viewpoint-relative content we manage to track persisting properties of objects. Thus within the foreground of experience, what Tye calls viewpoint-relative content is somehow related to viewpoint-independent content. We don't perceive an object as changing its properties, even though our experience does vary in its viewpoint-relative content.

Tye will no doubt respond that we don't perceive an object as changing its intrinsic properties because the latter are represented as staying the same by the viewpoint-independent layer of content. Yet we want to know how, given the changes in viewpoint-relative content that take place as we explore a scene, it is nevertheless possible for our experiences to have viewpoint-independent content that stays the same. Tye gives us no answer to this question.

It might be objected that by treating viewpoint-independent and viewpoint relative content as distinct layers of representational content, Tye's theory can avoid a problem that looms large for the sensorimotor theory. The perceiver has information about the intrinsic properties of an object on the basis of the viewpoint-independent content his experience carries. He also has information about these properties as seen from a particular location by virtue of the viewpoint-relative content his experience carries. However this can't be all there is to be said. Tye must surely concede that there is an interesting relationship between the viewpoint-relative and viewpoint-independent content. After all, the two layers of content do not represent entirely different properties. Consider again the two trees: the viewpoint-relative size my experience represents is one and the same property as the viewpoint-independent size. All that is different is that the viewpoint-relative layer of content represents the size of the tree from the perceiver's location. What is different is the way in which the property is represented. Once we recognise the active character of sensing, what is a problem for the sensorimotor theory also becomes a problem for Tye's theory. Some explanation must be given of how our experiences have not only varying viewpoint-relative content but also constant viewpoint-independent content.

The solution to this problem is to recognise that constancy is not something we can experience in a single take, but is rather something we experience over time. Consider what Pettit (2003) has called our abilities to "sift, sort and track" objects according to the properties they are represented as having. In

exercising these abilities we learn about or access an object's intrinsic properties. We distinguish one object from another and classify objects as being of the same or similar kinds on the basis of its perceived properties. We track them over time as they show up under different circumstance or move into and out of sight. We can in short identify "contrasts, commonalities and continuities" proper to objects we are seeing (Pettit 2003: 237). Pettit describes us as being empowered by our sensory encounters with the environment to make discriminations of this kind. However our discriminatory capacities extend beyond any single encounter with an object. We have expectations that were we to experience properties of the same type on different occasions or under different circumstances, we would be similarly empowered by our sensory encounters on each occasion. Forming expectations of this kind is, I claim, part of what it is to experience constancy. Notice however that these expectations of this kind depend on our sensorimotor knowledge. We expect the object to produce certain systematic changes in experience when it is presented under different conditions, or when we vary our sensory relation to it through movement. Our experience of constancy therefore depends on sensorimotor knowledge, but how?

I claim that as with the background we experience constancy across variations in experience by means of sensorimotor expectations. I agree with Tye that there is an aspect of the contents of experience that is viewpoint relative and that varies as our viewpoint changes. I have called this the foreground component of content in §3. However the foreground component is accompanied by sensorimotor expectations. These expectations concern how a property of a particular type will appear on future occasions as we vary our point of view on an object or if we were to vary the viewing conditions in which the object is presented. When my sensorimotor expectations are correct I succeed in tracking an object and its properties across variations in experience. To track a property across variations in experience – to identify a property as one and the same in different sensory encounters – just is for the property to exhibit constancy in our experience. We experience not only the way the object appears from a particular point of view now, but how it would appear as we vary our point of view. What we are experiencing is something that remains constant across variations in experience.

By acknowledging the contribution of sensorimotor expectations to the contents of experience we get an answer to question of how viewpoint relative

content relates to viewpoint independent content. Sensorimotor expectations are expectations relating to viewpoint relative content, and the ways in which an experience's viewpoint relative content varies as we change our point of view. Viewpoint independent content represents an object's intrinsic properties or what remains constant across variations in experience. I have just argued that an experience represents the latter by means of sensorimotor expectations. Viewpoint independent content is therefore built out of viewpoint relative content. An experience has viewpoint independent content by means of sensorimotor expectations.

So far I have been assuming that we are aware both of something remains constant in experience and of something that changes across different experiences. I have characterised the latter in terms of the way an object and its properties are presented from a location occupied by the perceiver's body. Thus in my discussion of colour constancy I said we experience the colour of the wall as being uniformly white but we are also aware of parts of the wall as having a different hue or brightness. Many have denied that we are aware of something that varies in experience as well as something that remains constant. They have doubted that when we experience the wall as uniformly white we are also aware of some parts of the wall as appearing lighter or darker than others. Sean Kelly (2008) has argued, for instance, that the way the different parts of the wall are presented in experience is something we can attend to, but it is not something we are ordinarily aware of. We become aware of this only be adopting a detached and reflective attitude towards our experience much as a painter might do in attempting to convey his experience of a scene on canvas. Schellenberg (2008: 78) has suggested that our awareness of these variations in appearance may form a part of the background to our experience. Unlike Kelly she accepts that they do form a part of what we ordinarily experience, but she agrees with Kelly that they are not something we normally attend to. Briscoe (2008) takes what is perhaps the hardest line on this issue in arguing that while we can notice the way an object appears we do so by engaging in an "exercise of visual imagination guided by seeing" (p.479). We are never literally aware of the ways in which objects appear. This is something we can bring ourselves to experience by engaging in acts of what Briscoe describes as "make-perceive".

I have been arguing that as we move our bodies fixating different parts of the visual field so the foreground of our experience changes. Whereas

Schellenberg holds that the foreground to our experience is occupied by intrinsic properties, I am claiming that the foreground is occupied by what we might call (following Noë) "perspectival properties". This is not to say that we are not also aware of an object's intrinsic properties. We are aware of an object's intrinsic properties when we form veridical sensorimotor expectations across distinct experiences. The properties that are visible to us at any given instant are perspectival properties. We see an object's intrinsic properties by means of the sensorimotor expectations that form a part of the content of every experience. It follows that an object's intrinsic properties have the same status as the background in our experience. They become visible to us through active exploration as we test out our sensorimotor expectations against the world. Until we engage in this active exploration an object's intrinsic properties form a part of the contents of experience only in the form of a set of sensorimotor expectations.

Why have so many philosophers denied experiencing anything that varies in different experiences of the same property? I suggest this denial is borne of a failure to recognise any difference within the contents of experience between background and foreground. If I am right, intrinsic properties show up in our experience only through the joint contribution of the background and foreground. A failure to recognise this distinction can make it seem like experiences only present us with intrinsic properties. A careful phenomenological description will, by contrast, reveal a foreground to experience that is constantly in flux, and a background to experience that ensures constancy.

Conclusion

I have defended a weak version of the sensorimotor theory, according to which sensorimotor knowledge accounts for two aspects of the contents of experience. Perceptual systems have mastery of sensorimotor contingencies. This mastery is what I've been calling sensorimotor knowledge. Sensorimotor knowledge is exercised in the formation of sensorimotor expectations. It is by means of these expectations that we have an experience of the background – the parts of a perceived scene we are not at that instant attending to. Sensorimotor expectations also give us an experience of constancy across variations in experience. When these expectations are veridical we experience an object's intrinsic properties despite always being presented with an object from a particular point of view.

In what remains of my conclusion I want to briefly return to Clark's argument against SSH in the light of what has just been claimed. Recall how Clark argued that conscious vision is vision for perception. Vision for perception is involved in the selection of targets for action and in the selection of types of action. The fine-grained control and guidance of action is under the purview of the vision for action system. Recall how the vision for perception system codes for objects and properties in space using allocentric coordinates. This means that what it represents are an object's intrinsic properties, not the properties of an object as they are coded relative to the perceiver's body. However I've argued that experience represents an object's intrinsic properties by means of sensorimotor expectations. These expectations are formed on the basis of sensorimotor knowledge. If my argument is correct it follows that experiences which are used for selecting types of actions and their targets, do not have contents that abstract away from the details of our sensorimotor engagement with the world. On the contrary it is by means of our sensorimotor knowledge that we are able to experience an object's intrinsic properties. Clark's argument against SSH fails.

I don't claim to have provided a full defence of SSH in this paper: I haven't shown that every difference in the contents of experience can be explained by sensorimotor knowledge. I have however attempted to show that sensorimotor knowledge is what explains two central aspects of the contents of experience. Moreover there is no incompatibility between the sensorimotor theory and two visual systems hypothesis. On the contrary, the vision for perception system depends on sensorimotor knowledge if it is to succeed in carrying out its functions.

References

Aglioti, S.M., Goodale, M., & DeSouza, J.F.X. 1995: "Size contrast illusions deceive the eye but not the hand." *Current Biology* 5: 679-85

Arend, L. & Reeves, A. 1986: "Simultaneous colour constancy." *Journal of the Optical Society of America A* 3.10: 1743-1751

Block, N. 2005: "Review of *Action in Perception*." *Journal of Philosophy* CII, 5: 259-72

Briscoe, R. 2008: "Vision, action and make-perceive." *Mind and Language* 23: 457-96

Brooks, R. 1991: "Intelligence without representation." *Artificial Intelligence* 47: 139-59

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

Clark, A. 2008: *Supersizing the Mind: Embodiment, Action and Cognitive Extension.* (Oxford: Oxford University Press)

Clark, A. 2006: "Vision as dance: three challenges for sensorimotor contingency theory." Psyche 12.1, downloaded from: http://psyche.cs.monash.edu.au/

Cohen, J. 2008: "Colour constancy as counterfactual." *Australasian Journal of Philosophy* 86.1: 61-92

Dennett, D.C. 1991: Consciousness *Explained.* (London: Penguin)

Ditchburn, R.W. & Ginsburg, B.L. 1952: "Vision with a stabilised retinal image." *Nature* 170, 36-37

Findlay, J. M. & Gilchrist, I. D. (2003). Active vision: The psychology of looking and seeing. (Oxford: Oxford University Press)

Gangopadhay, N. & Kiverstein, J. 2009: "Enactivism and the unity of perception and action." *Topoi* 28.1

Gibson, J.J. 1972: "A theory of direct perception." In J.R. Royce & W.W. Rozeboom (Ed.'s) *The psychology of knowing* (New York: Gordon & Breach), pp.215-40

Goodale, M. & Milner, D. 2004: *Sight unseen: an exploration of conscious and unconscious vision.* (Oxford: Oxford University Press)

Hurley, S.L. & Noë, A. 2007: "Can hunter-gatherers hear colours?" G. Brennan, R. Goodin, & M. Smith (Ed.'s) *Common minds: Essays in honour of Philip Pettit* (Oxford: Oxford University Press)

Hurley, S.L & Noë, A. 2003: "Neural plasticity and consciousness." *Biology and Philosophy* 18: 131-68

Hurley, S.L. 2001: "Perception and action: alternative views." *Synthese* 129, 3-40

Jacob, P. & Jeannerod, M. 2003: *Ways of seeing: the scope and limits of visual cognition.* (Oxford: Oxford University Press)

Kelly, S.D 2008: "Content and constancy: phenomenology, psychology and the content of perception." *Philosophy and Phenomenological Research* 76.3: 682-690

Kelly, S.D. 2004: "Seeing things in Merleau-Ponty." In T. Carman & M. Hauser (Ed.'s) *The Cambridge Companion to Merleau-Ponty.* (Cambridge: Cambridge University Press)

Krauskopf, J. 1963: "Effects of retinal image stabilisation on the appearance of heterochromatic targets." *Journal of Optical Society of America* 53: 741-44

Lycan, W. 1995: "Layered perceptual representations." E. Villanueva (Ed.) *Philosophical Issues* 7 (Atascadero: CA: Ridgeview), pp. 81-100

Mackay, D.M. 1962: "Theoretical models of space perception." In C.A. Muses (Ed.) *Aspects of the theory of artificial intelligence.* (New York: Plenum)

McConkie, G.W. & Rayner, K. 1979: "Is visual information integrated across successive fixations in reading?" *Perception and Psychophysics* 25: 221-24

McConkie, G.W. & Rayner, K. 1975: "The span of the effective stimulus during a fixation in reading." *Perception and Psychophysics* 17: 578-86

Merleau-Ponty, M. 1945/1962: *The Phenomenology of Perception* (translated by Colin Smith, London: Routledge)

Milner, D. & Goodale, M. 1995/2006: *The visual brain in action.* 2nd Ed. (Oxford: Oxford University Press)

Noë, A. (forthcoming): "Conscious reference." *Philosophical Quarterly*

Noë, A. 2008: "Reply to Campbell, Martin & Kelly." *Philosophy and Phenomenological Research* 76.3: 691-706

Noë, A. 2007: "Real presence." *Philosophical Topics* 33.1: 235-64

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

Noë, A. 2002: "On what we see." *Pacific Philosophical Quarterly* 83: 57-80

O'Regan, J.K. & Noë, A. 2001: "A sensorimotor approach to vision and visual consciousness." In *Behavioural and Brain Sciences* 24: 883-975

O'Regan, J.K. 1992: "Solving the "real" mysteries of visual perception: the world as an outside memory." *Canadian Journal of Psychology* 46.3: 461-88

Palmer, S.E. 1999: *Vision science: photons to phenomenology* (Cambridge, MA: MIT Press)

Peacocke, C. 1983: *Sense and Content: Experience, Thought and their Relations.* (Oxford: Oxford University Press)

Pettit, P. 2003: "Looks as powers." *Philosophical Issues* 13: 221-52

Philipona, D.L. & O'Regan, J.K. 2006: "Colour naming, unique hues, and hue cancellation predicted from singularities in reflection properties." *Vision Neuroscience* 23: 331-339

Riggs, L.A., Ratliff, F., Cornsweet, J.C., & Cornsweet, T.N. 1953: "The disappearance of steadily fixated visual objects." *Journal of the Optical Society of America.* 43: 495-501

Schellenberg, S. 2008: "The situation-dependency of perception." *The Journal of Philosophy* CV, 2: 55-84

Schwitzgebel, E. 2007: "Do you have constant tactile experience of your feet in your shoes? *Journal of Consciousness Studies* 14.3: 5-35

Siewert, C. 2005: Attention and sensorimotor intentionality. D. Woodruff-Smith & A.L. Thomasson (Eds.) *Phenomenology and philosophy of mind*. (Oxford: Oxford University Press), pp. 270-294

Snowdon, P. 1992: "How to interpret "direct perception"." T. Crane (Ed.) *The contents of experience.* (Cambridge: Cambridge University Press)

Tye, M. 2000: *Consciousness, Colour and Content.* (Cambridge, MA: MIT Press)

Thomas, N. 1999: "Are theory of imagery theories of imagination?" *Cognitive Science* 23.2: 207-45

Thompson, E. 2007: *Mind in Life:* (Cambridge, MA: Harvard University Press)

Weatherson, B. 2006: "Intrinsic vs Extrinsic Properties." In E. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*

Wollheim, R. 1968/1980: *Art and its objects,* 2nd Ed. (Cambridge: Cambridge University Press)

Yarbus, A.L. 1967: *Eye movements and vision.* (New York: Plenum)