ALVA NOË

EXPERIENCE AND THE ACTIVE MIND

ABSTRACT. This paper investigates a new species of skeptical reasoning about visual experience that takes its start from developments in perceptual science (especially recent work on change blindness and inattentional blindness). According to this skepticism, the impression of visual awareness of the environment in full detail and high resolution is illusory. I argue that the new skepticism depends on misguided assumptions about the character of perceptual experience, about whether perceptual experiences are 'internal' states, and about how best to understand the relationship between a person's or animal's perceptual capacities and the brain-level or neural processes on which they depend. I propose a conception of perceptual experience as a form of skillful engagement with the environment on the part of the whole person or animal.

> We find certain things about seeing puzzling, because we do not find the whole business of seeing puzzling enough.
> – Wittgenstein (1953/1958, 212)

1. INTRODUCTION

My topic in this paper is a species of skeptical reasoning about visual experience that takes its start from developments in perceptual science. Traditional skepticism about experience questions whether we can know, on the basis of experience, that things are as they are perceptually experienced as being. The new skepticism presses instead the question whether in fact we have the perceptual experiences we think we have. According to this skepticism, the impression of visual awareness of the environment in full detail and high resolution is illusory. Perceptual consciousness is, in a fundamental sense, false consciousness. The purpose of this paper is to challenge this new skepticism. I argue that it rests on misguided assumptions about the character of perceptual experience, about whether perceptual experiences are 'internal' states, and about how best to understand the relationship between a person's (or animal's) perceptual capacities and the neural processes on which they depend.

Synthese **129:** 41–60, 2001.
© 2001 *Kluwer Academic Publishers. Printed in the Netherlands.*

2. THE NEW PUZZLE ABOUT VISUAL CONSCIOUSNESS

To ordinary reflection, visual experience seems to represent the visual world in sharp focus and rich detail. This thought is expressed by Gregory, who writes:

> We have only to open our eyes, and spread before us lies a banquet of colours and shapes, shadows and textures: a pageant of rewarding and threatening objects, miraculously captured by sight. All this, from two tiny distorted upside-down patterns of light in the eyes. (Gregory 1966/1978, 1)

The idea is that when we see we are aware, all at once, of the whole, variegated scene. Vision seems to allow us to take in the array of objects at a glance. This would seem to be a striking feature of ordinary perceptual consciousness.

Visual scientists have long sought to explain how we can enjoy visual impressions of such richness given the limited information made available to the brain in the form of patterns of light on the retina. In addition to the facts noted by Gregory (that the retinal image is distorted, tiny, upside-down), there is the fact that the resolving power of the eye is limited and non-uniform. Outside the high-resolution central (foveal) region, the retina is nearly color-blind and its powers of discrimination are severely limited. In addition, the eye is in nearly constant motion, saccading from point to point in the visual field three or four times a second. During saccades there is a total shutdown of new retinal processing. As a result of this saccadic suppression, the data made available to the retina takes the form of a succession of alternating snapshots and gray outs. How, on the basis of this fragmented and discontinuous information, are we able to enjoy the impression of seamless consciousness of the environment in all its detail and complexity?

Traditional visual science would have it that information contained in successive fixations is integrated to produce a stable, detailed brain-model or representation corresponding to what we in fact perceive – to the impression of the whole scene (Feldman 1985; Trehub 1991; both discussed in Rensink 2000). On this approach, vision is a process whereby the patchy and fragmentary representations of the earliest stages of vision are transformed into the detailed, stable representations underlying actual perceptual experience.[1]

Recent work in perceptual psychology challenges this way of thinking of vision by challenging the conception of visual experience that seems to be presupposed by it. Work on "change blindness", for example, demonstrates that we tend only to notice changes to the visible environment if the changed feature affects the gist or meaning of the scene (Simons and Levin

1997), or if we happened to attend to the change when it occurs (Rensink et al. 1997). Recent experiments show that normal perceivers are frequently unable to detect changes to important, visible, even salient features of visual scenes. (For a detailed review of the change blindness literature, see Intraub 1997; Simons and Levin 1997; Simons 2000.) The experiments proceed by presenting changes to perceivers while at the same time disrupting normal attention-directing mechanisms. For example, flashing "mud splashes" on the monitor when changes occur in photographs of natural scenes swamps the motion detectors which would otherwise direct attention to the location of change, and so prevents us from noticing such changes as whether a traffic line is broken or unbroken, or whether the cathedral in the center of the field of view is displaced substantially to the left (O'Regan et al. 1999). Importantly, despite our impression of awareness of the scene in its totality, we seem actually to be conscious of only a limited amount of what we think we see.

The same conclusion is thought to follow from the work of Mack and Rock on "inattentional blindness" (Mack and Rock 1998; see Noë and O'Regan 2000, for discussion of this phenomenon). They demonstrate, in striking and surprising ways, that perceivers only "see" that to which their attention is directed. For example, you will not notice a person in a gorilla-suit walking through a basketball game if you are focusing your attention on the number of times one team gains possession of the ball (Simons and Chabris, in press). The impression of seeing all the detail at once, then, however robust, is an illusion, or so these results seem to show.[2]

Consider a second example of the skeptical reasoning at work. As noted above, the resolving power of the retina is severely restricted at its periphery, both with respect to spatial detail and color. As noted by Dennett (1991, 53), you cannot tell the color of a playing card held at arm's length at the periphery of the visual field. This simple experiment shows, or so Dennett (1991, 53) suggests, that "naïve reflection" is deceived as to the character of the visual field. We are the victims of an illusion when we have the impression that we experience color-detail distributed uniformly across the visual field from the center out to the periphery.

Dennett makes a similar claim in connection with his influential critical discussion of "filling-in" (1991, 1992). It is well-known that there is a region of the retina where there are no light sensitive receptors. This "blind spot" notwithstanding, we do not have the impression that there is a gap, or hole, in the visual field. In fact, it seems – "to naïve reflection" – that we are aware of the region of the visual field corresponding to the blind spot. It seems to us, that is, as if the brain has filled in the gap. Dennett challenges the inference from this observation to the proposition that there

must indeed occur filling-in to make up for the absence of information at the blind spot. He observes that the effect in experience would be the same if the brain, instead of actively filling in, simply ignored the absence of information about the region corresponding to the blind spot. (For a detailed discussion of filling-in and Dennett's position, see Pessoa et al. 1998a,b.)

As far as we can tell on the basis of experience, then, the brain might not fill in at all. But if it should turn out that the brain does not fill in, Dennett reasons, then we are mistaken in our conviction that we actually undergo a perceptual experience as of an unbroken expanse. We could not experience such an expanse if there is no filling in. Note that the mistake here, the illusion (if such it be), is not a perceptual one to the effect that there is an unbroken expanse of wall. The illusion, rather, pertains to whether or not we actually have an experience as of an unbroken expanse. The illusion is one of consciousness.

The upshot of this critical reflection is to call into question whether we (the theorist and lay person alike) are right to believe, as it is supposed we do, that when we see we enjoy a detailed picture of the whole visual field. As far as consciousness is concerned, things are not as they seem! In so far as traditional visual theory rests on such a conception of experience – vision as the process whereby the brain constructs the detailed internal model needed to produce the detailed experience we take ourselves to enjoy – the new skepticism is a challenge to the guiding ideas of traditional visual theory.

3. VIRTUAL REPRESENTATION

For traditional approaches to vision, the central fact to be explained is that perceptual consciousness greatly outstrips what is given to the visual system in the form of irradiation of the retina. If the new skepticism is right, what cries out for explanation is not how experience can be so rich given such impoverished initial data, but rather, how it is that it *seems to us* as if experience is so rich when in fact it is not. Dennett writes: "One of the most striking features of consciousness is its *dis*continuity. Another is its apparent continuity" (1992, 48). Against the background of these considerations, the problem that emerges for visual theory is to explain the apparent continuity of perceptual experience despite its actual discontinuity.[3]

O'Regan (1992), O'Regan et al. (1999) and Rensink (2000) have suggested that the solution to this problem lies in something like what Rensink calls "virtual representation" (Rensink 2000). The visual system does not

produce detailed representations of the environment, as the traditional perspective assumed – if it did, then we would not suffer from change blindness – for all intents and purposes, however, it is *as if* it does. Rensink (2000) explains the proposal like so:[4]

> Why do we feel that somewhere in our brain is a complete, coherent representation of the entire scene? .... The explanation for this may lie in the dynamic nature of the representations involved .... This realization gives rise to the idea of *virtual representations*: instead of forming a coherent representation for all objects, represent only those objects needed for the immediate task at hand. If memory and attention can be coordinated so that a coherent representation of an object can be formed whenever required, the representation of the scene will appear to higher levels as if it is "real", i.e., as if all the objects it represents are present simultaneously. Such a representation will then have all the power of a real one, while requiring far less in the way of processing and memory resources.

The idea behind this proposal is that the visual system does not need to produce and store a detailed representation of the environment, since it has speedy and reliable access to information in the environment when the information is needed. A system that gathers information reliably and on-time will be equivalent, with respect to the availability and usability of information, to one that stores representations. In other respects, however, it will be superior. Consider that the virtual system does not need to bear the burden of producing the detailed internal representations in the first place.

The proposal that vision deploys virtual representations has broad implications for visual theory and the philosophy of mind. (1) Whereas the traditional view supposes that visual computation occurs only in the retina-brain system, according to the new approach, computation depends also on the non-brain body. For example, the virtual representation mechanism exploits our capacities for eye movement. (2) Whereas traditional constructive views emphasize processes going on inside the animal in the retina-brain system, the new VR model (virtual representation model) depends on the fact that the perceiver is situated in and capable of acting on a stable environment.

This VR proposal thus dovetails nicely with new cognitive science research programs that increasingly reject the assumptions of traditional visual theory in favor of 'active' models. If vision is not a process of building internal models of the outer scene, then what is its function? One obvious answer to this question, as Milner and Goodale (1995) point out, is that vision enables the performance of skilled action such as walking through a crowd and catching a ball, in addition to enabling perceivers to learn about the structure and layout of the environment. The 'animate vision' program of Ballard (1991), and the 'embodied cognition' program of Brooks (1991) take as their starting point the assumption that

the primary function of vision is to enable action-guidance and they leave it as an open question, to be settled on the basis of empirical research, whether such action-guidance must depend on the construction of detailed representations. Brooks proposes that we model the internal organization of intelligent systems by means of what he calls a "subsumption architecture". The idea is that we view the creature as constructed out of the interaction of more or less autonomous activity-producing subsystems (Brooks 1991; Haugeland 1995; Clark 1997). A subsumption architecture is a system of horizontally rather than hierarchically organized modules, each of which is, as Hurley puts it, "a complete input-output-input loop, essentially continuous and dynamic, involving external as well as internal feedback ...". In this way, "not only are sensory and motor processes coupled, but the neural network is coupled to the creature's environment ..." (Hurley 1998, 408). The significance of the change blindness and inattentional blindness results is best appreciated against the background of the emergence of these anti-constructivist theories. Phenomena such as change blindness demonstrate that the brain does not produce models of the detailed environment to support vision, and this finding provides support for the newer anti-constructivist approaches.

4. VIRTUAL REPRESENTATION AND INTERNALISM

Most visual scientists assume that for every perceptual experience, there is a neural substrate that is sufficient, as a matter of scientific law, to produce it (van Gulick 1998).[5] Teller and Pugh call such an immediate neural substrate "the bridge locus" (Teller and Pugh 1983). Crick and Koch give expression to the same "internalist" assumption in their talk of "the neural correlates of visual awareness" (Crick and Koch 1990, 1995, 1998).

Let us say that a mental state is internal if neural or bodily states are sufficient as a matter of scientific law to produce that state. Traditional theories of vision represent perceptual states as internal in this sense. As we have noticed, the central aim of visual theory has been to explain how the internal neural substrate of experience is produced on the basis of initial retinal input. A state is external, in contrast, if neural states and events are not sufficient as a matter of law to produce it. If Putnam (1975a) and Burge (1979) are right, then beliefs are in this sense external states. Beliefs are at least in part constituted by body-external events, processes and conditions. Recently Clark and Chalmers (1998) have argued that certain kinds of cognitive states – states of, e.g., thinking, calculating, navigating – may be at least partially external in that, at least sometimes, these states depend on making use of symbols and artifacts that are outside the human body.

This develops a point first made by Wittgenstein (1958),[6] but which has received a good deal of attention in the recent work of roboticists and cognitive scientists.

Against this background of philosophical externalism about belief and other states, perceptual states present particular difficulties. In so far as they are phenomenal states, like sensations, they would seem to be internal (in our sense). But in so far as they are intentional states, like beliefs, they would seem to be world-involving.

One might think that that the active vision framework provides a challenge to the internalist framework of traditional visual theory. (See Pessoa et al. (1998a,b) and Noë et al. (2000) for an argument to this effect.) Whereas traditional theory holds that vision is a process in the retina-brain computational system, the active vision paradigm suggests that visual processing in fact involves the non-brain body and the environment. For the active approach, it would seem, no neural substrate could be sufficient, as a matter of law, to produce visual experience. The point can be made by means of a comparison. One might reasonably believe that materialism would guarantee that the engine of the automobile is sufficient, as a matter of law, to produce the driving action of the car. But this is clearly not so. No matter what the engine does, a car without wheels will not drive. Nor will a car drive if it is upside down, or suspended from a hook, or placed in a swamp. The mechanical substrate is sufficient only given the embodiment of that substrate in a normal vehicle and the appropriate embedding of that vehicle in a normal environment.

For analogous reasons, neural substrates, it is tempting to argue, would not be sufficient, as a matter of law, to produce perceptual states, unless further conditions are satisfied, viz. conditions of embodiment and situatedness. According to the VR approach, visual experiences will be produced by the interaction between the nervous system, the non-brain body and the environment. As a result, perceptual states would depend on brain/body/environment complexes, but not on neural states on their own. In the sense of the terminology we have adopted here, this would make perceptual experiences external states. (Notice that the claim that the neural substrate is not sufficient, on its own, to produce perceptual experience does not rest on any anti-materialistic commitments (*pace* van Gulick 1998, 777, cited above.)

It now seems misleading, however, to characterize the VR and activity-based approaches to vision as externalist. For such views leave unchallenged the basic idea of internalism: that to see detail, the detail must be represented in the head. They provide an account of the workings of the visual system that gives important place to the non-brain body and

the environment, to be sure. But it remains crucial, for such models, that it is only in a *virtual* sense that we visually experience the detailed environment. The VR approach, that is, takes for granted that we do not *really* see all the detail, since, after all, the detail is not in the head. It then provides an account – appealing to facts about embodiment and environmental situatedness – to explain why it seems to us (erroneously) as if we see (have in the head) so much more detail than we really do. This is nicely expressed by O'Regan, who writes: "despite the poor quality of the visual apparatus, we have the subjective impression of great richness and "presence" of the visual world: But this richness and presence are actually an illusion, created by 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, possibly rendered available by an eye movement" (1992, 484).[7]

Active vision theories purchase their would-be externalism at the price of abandoning perceptual experience as mere confabulation. A genuine externalism about perceptual experience would insist that the absence of corresponding internal representations does not impugn the "reality" of experience. Experience of detail, such an externalism would argue, does not require that the detail be modeled in the head. This more thoroughgoing externalism would reject the "models in the head" conception that characterizes traditional visual theory *without* embracing the new skepticism.[8]

5. RECOVERING PERCEPTUAL EXPERIENCE

Does it really seem to us (to lay people as well as philosophers and scientists) that when we see we have all the seen detail in our consciousness all at once, and that what is seen is in uniformly sharp focus from the center of the visual field out to the periphery, as both the traditional theorist and the new skeptic insist?

Let's revisit the skeptical reasoning in the context of a simple example from touch. You hold a bottle in your hand with eyes shut. You make contact with the bottle at isolated points where the bottle touches your skin. Nevertheless, it seems to you as if the whole bottle is present to your awareness, as if you perceive the whole bottle. The traditional view would explain this by supposing that the brain produces an internal model on the basis of information made available in one's finger-to-bottle points of contact and that it is this internal construction which enables us to have the experience of the bottle we take ourselves to have. In rebuttal the new skeptic charges that the whole exercise rests on a false characterization of

what the experience is like. Contrary to the way it seems, we do not really experience the untouched parts of the bottle. The experience of all that detail is a confabulation.

But do we really subscribe to the questionable conception of what our experience is like? You surely take yourself to be aware of the bottle, to sense it as a whole and as present. But it seems far-fetched to say that it seems to you as if you are in contact with all parts of the bottle's surface. Rather, it seems to you as if the bottle, as a whole, is present, and that you have access to detail about currently untouched parts of the bottle by further hand movements. Just the same sort of point can be made about vision. When we see, it perceptually seems to us as if we encounter a world of immense detail, but it doesn't seem to us as if we have all that environmental detail in our minds at once, no more than it seems to us as if we make contact with all parts of the bottle's surface. Rather, it seems to us as if the environment is there, and as if we can extract information from the environment by moving our eyes or head or by reorienting our bodies.

A very similar distortion of our thoughts about our perceptual experience is in evidence in connection with Dennett's discussion of the blind spot (as noticed by Pessoa et al. 1998a,b). It is outrageous to suppose that it seems to us, naïve perceivers, as if there is filling-in at the blind spot. It is true that we take ourselves, even with one eye shut and our gaze fixed straight ahead, to be aware of the uniform expanse of the red wall. But it is not the case that we take ourselves to have the whole of that expanse in view at once. And so it is not the case that we must account for our supposed impression that we have the whole expanse in consciousness at once. This in turn is compatible with its seeming to us, when we see, as if that which we see – in this case, the wall – is present all at once for viewing. In addition, it will be no part of our impression that it seems to us as if the brain fills in the region corresponding to the blind spot. First, thanks to eye movement, there is no one region corresponding to the blind spot. Second, the very existence of a blind spot, or of any putative processes of neural filling-in, remain entirely below the threshold of consciousness. Whatever methods are employed by the brain and nervous system to enable us to see, we take ourselves to enjoy the appearance of the unbroken wall. And we do.

Exactly similar points can be made about the change blindness and inattentional blindness studies. Rensink asks (see quotation above): "Why do we feel that somewhere in our brain is a complete, coherent representation of the entire scene?" But the question rests on a false presupposition. It does *not* seem to us as if somewhere in our brain there is a complete, coherent representation of the entire scene. It is not part of our experience

or our pretheoretical thought about experience that we construct models in the brain. As stressed by Pessoa, Thompson and Noë, perceptual experience is directed to the world, not to the brain (Pessoa et al. 1998a,b; Noë et al. 2000). To suggest that it seems to us when we see as if the brain produces a detailed model of the environment is to falsify our experience. It does seem to us, as O'Regan writes, that we have the impression of the complete visual world outside of us. But we do! We take ourselves to be confronted with and embedded within the detailed, high-resolution environment. The fact that it is not possible to enjoy that awareness on the basis of an individual fixation is irrelevant. For our awareness of the complete and detailed environment is not an awareness of ourselves as representing that environment to ourselves in the form of snapshot-like fixations, as it were all at once. It is an awareness of the environment.

It is no doubt correct that our ability actively to probe the environment, making use of eye and head movements (for example) is crucial to understanding the character of perceptual consciousness, as O'Regan (1992) has emphasized. But it is a mistake to cast this (very important) discovery in terms of the sustaining of an illusion to the effect that we have the detail in consciousness at once. We do not have all the detail in consciousness at once. Nor does it seem to us as if we do. It seems to us, rather, as if the environmental detail is present and as if we are ourselves embedded in that environment. Although the skeptic is correct that the "detail in the head" conception of perceptual consciousness is misguided, it is a mistake to suppose, as both the skeptic and the traditional theorist seem to do, that we (ordinary folk) subscribe to the "detail in the head" conception.

Consider once again the example of feeling the bottle. Notice that the traditional theorist's supposition that the basis for our awareness of the bottle is the availability of an internal representation is, in effect, an unnecessary shuffle. As noticed by O'Regan (1992), there's no need for an internal model of the bottle, since the bottle itself is right there, in our hands, to serve (in Brooks's (1991) phrase) "as its own best model". Our perceptually-guided exploration of the bottle does not require that we have access to a mental model of the bottle, since the bottle itself, thanks to its real shape and tangibility, guides our movements and our explorations.

This suggests the possibility of a way of thinking about perceptual experience as, precisely, a form of active engagement with the environment. In the case of the bottle, it is not our possession of an internal model of the bottle that is the basis of our contact with it – indeed, we do not have that sort of contact. Rather, it is the fact that we are so related to the bottle that we are, as it were, ready and able to acquire information as need arises. And so in the visual case: the ready availability of environmental detail,

and – this is very important – the skill-based confidence on the part of the perceiver that he or she is able to acquire that detail through movement, is the basis of our feeling of the presence of the environment as a whole.

The ground, then, of the feeling of perceptual contact with the whole detailed environment, whether in sight, touch, or other perceptual modalities, is the possession by the perceiver of skills for exploration (O'Regan and Noë, under review). When you see a cat through a picket fence, you see the whole cat even though, in a sense, you only see the parts that are visible between the slats of the fence. It is sometimes supposed that we infer the presence of the unseen parts of the cat. What is striking is that, phenomenologically speaking, one experiences the unseen parts of the cat as perceptually present and not merely as an inferred presence. To bring this out, contrast the sense in which we take the unseen parts of the cat to be present with the sense in which we take, say, the hallway outside our room to be present. There can be no doubt that the hallway seems present, but it does not seem to be perceptually present. The present proposal attempts to account for this difference. The basis for this "feeling of perceptual presence" is one's skill-based confidence in the effects of probing on what is seen, e.g., if you step to the right a bit, formerly unseen bits of the cat come into view; if the cat moves, new cat surface becomes visible. The feeling of presence resides in the immediate accessibility, through control of one's sense organs, of detail that is present there all along.

On this conception, perceptual experiences are things we do, not events in the head. Visual experience, like the tactual awareness of the bottle, is a temporally extended mode of exploratory contact (e.g., with a bottle). Seeing is in an activity of exploration as mediated by the characteristic sensorimotor contingencies of vision (e.g., by such facts as that, when we blink, we don't see; that when we move the eyes to the right, the retinal projection of an object in the environment moves to the left; that when we move around a partially occluded object, hidden surface comes into view). To be a perceiver, on this view, is to be familiar with and skilled at coping with the myriad ways in which how things look (or feel) depends on what one does.

Perceptual awareness, on this view, is a state of interactive engagement with the world, not a state of picture-making. It is just such an active (or "enactive") conception of perceptual consciousness that is endorsed by reflection on the character of perceptual experience.[9] Perceptual consciousness is a capacity of the whole, environmentally situated person or animal.[10]

6. THE PRIMACY OF THE PERSONAL AND THE REFUTATION OF INTERNALISM

Once we reject the dogma that in normal thought and life we take ourselves to have all environmental detail in consciousness at once, we are in position to rethink the internalist notion that neural states must be sufficient to produce experience. Indeed we may want to ask: how *could* it be that visual experience is produced by neural activity? The worry is not that there is an explanatory gap. The worry is that, given the temporally extended, whole-animal character of perceptual activity, it becomes unclear how one could even come to think that neural activity could be nomically sufficient to produce experience.

A somewhat simplistic comparison may help. For the sake of argument, let us assume (contrary to fact) that no action is taken on the part of the US government without a deliberate decision-making act on the part of the President. The fact that an action on the part of the President is in this way necessary for all governmental undertakings is compatible with its being the case that the decision-making acts on the part of the President are not sufficient to produce the corresponding action on the part of the government. Indeed, this is exactly what we should expect. The President's power resides not in his person, but in his office, and only thanks to the larger context provided by the Constitution. It is thus only on the assumption that a broad range of further conditions are met that the President's personal decision making is effective. In addition, as a practical matter, the President is influenced by, and must answer to, political and social pressures. Sometimes the President's hand is forced. As a consequence of all this, it is ludicrous to think that you could explain all governmental action with reference to decision-making acts on the part of the President alone.

If we accept that perceptual experience is a mode of active engagement with the environment, as I have proposed, then we should accept that neural states are *presidential*, in the sense of the previous paragraph. Neural states, that is, contribute to the production of experience only given their embodiment in the whole living body (which in turn must be embedded in an environmental setting). There is no more reason to believe that a neural state could be sufficient to produce the experience of seeing than there is to think that the President's decision-making acts are sufficient to produce governmental action.

The analogy with the powers of the Presidency breaks down when we consider that whereas the President is not powerful enough to produce governmental action all on his own – after all, the powers of the Presidency

are not intrinsic to the man who is President – perceptual experiences are just not the sort of thing that can be caused by patterns of neural activity (no matter how powerful). For perceptual experience is not just something that happens, all at once, like a headache. It is something that is performed – enacted – by a living animal engaged in a pattern of exploration of its world.

The upshot of this argument is not that neural states do not produce experience. The upshot rather is that they do so only given a great deal of stage setting (embodiment, environmental situatedness, etc). As a consequence of facts such as these, it is a mistake to suppose – as van Gulick and Crick and Koch seem to – that perceptual consciousness can be understood at the brain-level.

Perceptual consciousness, I conclude from this, is not an internal, subpersonal phenomenon. It is a phenomenon not of the brain but rather of the whole animal. Investigations into the neural basis of perception must heed this fact at the risk of producing simplistic and ultimately unpromising models. We have already encountered one example of the sort of error to which one committed to a brain's eye view will be prone: the assumption, evident in many of the writers canvassed here, that the absence of a neural substrate sufficient to produce experience is a reason to doubt that we really have the sort of experience in question. (See Pessoa et al. 1998a,b for more on this.) In closing, let us briefly consider a different sort of phenomenon.

Research on prosthetic vision by Bach-y-Rita et al. (1969) and Bach-y-Rita (1996) provides a concrete example of the autonomy of perceptual consciousness with respect to the brain level. In tactile-visual substitution systems (TVSS), optical images picked up by a camera (worn, say, on the head) are transduced in such a way as to activate an array of stimulators (vibrators or electrodes) in contact with the skin (on, e.g., the abdomen, back or thigh). Optical images in this way produce a localized pattern of tactile sensation. After an initial period of training, congenitally blind subjects cease to experience tactile sensations when they use the TVSS device, and come to report that they experience objects as arrayed out before them in three-dimensional space, just as captured by the camera. Such tactile-perception enables subjects to make judgments of shape, size and number and also to perceive spatial relationships between things, of the sort normally made by vision. With sufficient practice, subjects are able to engage in tasks requiring skillful sensorimotor coordination, e.g., batting a ball or working on an assembly line. In light of this, we might be justified in concluding that Bach-y-Rita's TVSS enables subjects to see, that tactile-vision is indeed a form of vision.

Let us question: what would this conclusion that TVSS is a form of vision amount to? This much is clear I think: despite the superficial fact that tactile-visual information does not come to us through the eyes, it seems to present the environment in a manner structurally like that of vision. The resulting perceptual experiences are clearly non-tactual; objects are experienced as arrayed out in space before one, and not, as coming into contact with one's skin. Moreover, changes in one's position and orientation relative to the environment yield changed glimpses of the scene in manner that is structurally like that of vision. This is evident in the applicability, to tactile vision, of such notions as that of perspective, occlusion, change in apparent size, etc. Furthermore, the manner in which the prosthesis is integrated into perceptual activity is vision-like: to see something, you direct your camera to it; to see it better, you move the camera closer; to see it despite an occluding obstacle, one must move around the obstacle. But there is an even stronger sense in which tactile-vision depends on the successful integration of the prosthetic apparatus into a *vision-like* sensorimotor scheme. A subject's initial use of TVSS yields only tactile sensations – the effect of the vibrators on the skin. It is only when the subject begins to integrate those sensations into patterns of responsive movement (appropriate responses to stimulation) that the tactile sensations lose their tactile quality and become (or seem to become) visual.[11]

The question, *Is tactile vision truly a form of vision?*, is not, at least in the first instance, a question about the character of the neural activity that is necessary for its occurrence. Note that the neural basis for TVSS is (presumably) located in the somatosensory cortex. If we decide that TVSS is a form of vision, must we conclude that there must be neural activity in visual cortex in addition to whatever is going on in somatosensory cortex? It is an open question whether tactile vision produces activity in visual cortex. But it seems clear that we are not entitled to believe that the *visual* character of tactile vision alone entails that there must be such activity. What is clear is that we can perfectly well make sense of the *possibility* that the forms of sensorimotor integration characteristic of tactile vision could occur in the absence of activity in visual cortex. The *visual* character of these experiences resides in the activity of the whole integrated animal. What this shows is that vision is not a neural phenomenon. It is a phenomenon at the level of the whole animal.

7. CONCLUSION

In this paper I have proposed a way of thinking of perceptual experience as a temporally extended form of skillful engagement with the world. (See O'Regan and Noë, under review, for extended development of this idea.) From the standpoint of this new conception it seems clear that perceptual consciousness is not, as numerous authors have claimed, a "grand illusion". When it comes to perception, we are not all theorists who are committed to the "details in the head" conception, as Dennett (1998) has suggested. We take ourselves, in experience, to be aware of the densely detailed environment, but we don't take ourselves to have all that detail in consciousness at once. Rather, we take ourselves to have access to that detail. I argue that our sense of the presence of the detail consists in our skill-based confidence in our ability to acquire the information by movements of the body (as argued in O'Regan and Noë, under review). Finally, against the background of this conception of experience as a form of skillful engagement, it seems unlikely that we can explain the role of the brain in perception and perceptual experience as, so to speak, the seat of consciousness and as sufficient to produce experience. To understand the role of the neural in perception we must have clearly in view the fact that perception is a phenomenon of the whole animal.

ACKNOWLEDGEMENTS

This paper grows out of collaboration with Evan Thompson and Luiz Pessoa, and also with Kevin O'Regan. I owe them a debt of gratitude. Thanks also to Daniel Dennett, whose work provides the starting point, and also to David Chalmers and Sean Kelly, for very helpful discussion. I benefited greatly from discussion of this material at Erik Myin's Cognitive Science Conference on Perception, Consciousness and Art at the Vrije Universiteit Brussel in May 1999. I am especially grateful for the comments received there from Jennifer Church, Fred Hamker, Susan Hurley, Marcel Kinsbourne, Erik Myin and Joëlle Proust. I would also like gratefully to acknowledge the support of faculty research funds granted by the University of California, Santa Cruz.

NOTES

[1] According to this conception of vision, vision is a process of "inverse optics". The goal of vision is to recover a description of the 3D environment, from the retinal image. Marr

(1982, 2) wrote: "vision is the *process* of discovering from images what is present in the world and where it is". The thought is expressed well by Tsotos (1987), who says "The goal of an image-understanding system is to transform two-dimensional data into a description of the three-dimensional spatio-temporal world" (p. 389) (cited in Churchland et al. 1994). Marr (1982) took his characterization of vision as a process whereby the brain produces descriptions (or models) of the environment to be a formulation of what vision is at the "computational level of analysis". Such a formulation is preliminary to investigations in to the ways in which that computational problem can be solved by an actual brain. Milner and Goodale sum up the position like so: "It is commonly assumed that vision in humans has a single function: to provide a unified internal representation of the external world which can then serve as the perceptual foundation for visually based thought and action. From this perspective, the task of visual science boils down to the problem of understanding how the spatiotemporal mosaic of light striking the retina is parsed into the array of discrete objects and events that comprise one's perceptual experience of the world" (1995, 5).

[2] Many psychologists and philosophers are committed to the idea of this "grand illusion hypothesis". See, for example, discussion of Dennett in the text below, and also Ballard (1991, 1996), Dennett (1991, 1992, 1998), Churchland et al. (1994), Blackmore et al. (1995), Rensink (1999). For a general overview, see Noë et al. (2000). Blackmore and her colleagues (1995, 1075) write: "we believe that we see a complete, dynamic picture of a stable, uniformly detailed, and colourful world", but "[o]ur stable visual world may be constructed out of a brief retinal image and a very sketchy, higher-level representation along with a pop-out mechanism to redirect attention. The richness of our visual world is, to this extent, an illusion". This citation nicely illustrates the skeptical reasoning.

[3] This critical question is posed by Ballard, and he also provides the now familiar skeptical answer (1991, 59–60; see also 1996): "Saccades at the rate of three per second are the routine in visual problem solving. Furthermore most of the brain structures that represent visual information are retinally indexed. This means that their state is changed with each eye movement. This raises a technical puzzle for human visual perception: How can the world appear to be stable when the data collecting process is so dynamic? We believe that this is a profound question with a surprising answer: The visual system provides the illusion of three-dimensional stability by virtue of being able to execute fast behaviors. This point may be very difficult as it is so counter-intuitive . . .". The task of visual theory, then, is precisely to explain the mechanisms by which the illusion of perceptual consciousness of the stable, detailed environment is produced.

[4] O'Regan et al. (1999, 34) write: "If only attended parts of the environment are represented in the brain, how can we have the impression of such richness and completeness in the visual world outside us? The answer might be that the visual world acts as an external memory. We have the impression of simultaneously seeing everything, because any portion of the visual field that awakens our interest is immediately available for scrutiny through the unconscious flick of the eye or of attention".

[5] Van Gulick writes (1998, 777): "Assuming the truth of materialism . . . every token percept must have a neural substrate – whether local or global – that is sufficient to produce or constitute it".

[6] "If again we talk about the locality where thinking takes place we have a right to say that this locality is the paper on which we write or the mouth which speaks" (Wittgenstein 1958, 7).

[7] The same thought is expressed by Churchland, who writes: "Although unattended objects may be represented in some minimal fashion (sufficient to guide attention shifts and

eye movements, for example) they are not literally seen in the sense of "visually experienced" (pp. 26–7). O'Regan, it should be mentioned, now acknowledges the error in talk of a "grand illusion".

[8] Notice that the account of internalism offered here obliterates the familiar distinction between internalism about content and internalism about vehicles of content (see Hurley 1998 for detailed discussion). According to this distinction, Putnam and Burge are content externalists. They believe that content depends on relations between the individual and his or her environment *whatever* the character of the vehicles of content. Most content externalists are also vehicle internalists. That is, even though they hold that content is fixed by facts about environmental embedding, they insist, nevertheless, that what carries content (e.g., states of the brain) must be internal. Putnam and Burge are probably both vehicle internalists in this sense. The externalism of Clark and Chalmers, on the other hand, is a species of vehicle externalism. They argue that artifacts external to the body (e.g., a map) are part of that which carries content. This still leaves it an open question whether they are also content externalists. One can be a vehicle externalist and still hold that content is not constituted by facts of environmental embedding. As I use the terms, content externalism and vehicle externalism are each species of externalism. To be an externalist, in my sense, entails that one is either a content externalist or a vehicle externalist (or both). Content and vehicle internalism, likewise, are both species of internalism in my sense. But to be an internalist, in my sense, requires commitment to both content and vehicle internalism, since, as we have seen, to give up either of these commitments would be to accept that what is going on in the head is not sufficient to produce experience.

[9] The term "enactive" is used by Varela et al. (1995) to indicate a way of thinking about the mind according to which (1) the subject of mental states is taken to be the embodied, environmentally situated animal; (2) the animal and the environment are thought of as a pair, standing in a relation of being essentially coupled and reciprocally determining; (3) perceptual and other cognitive states are thought of in terms of activity on the part of the animal and as nonrepresentational; (4) the mental life of a creature is taken to be an autonomous domain for the sort of investigation pursued within the philosophical movement known as Phenomenology. In the text I characterize what I call the active conception of experience as a way of thinking about experience according to which (1) experience is a phenomenon of the personal level, (2) experience is to be understood in terms of the interactive engagement of the animal with its environment, and (3) experience is, in the sense explained above, autonomous with respect to the subpersonal system.

[10] Patricia Churchland has criticized the active conception of experience (in discussion) by holding (1) that we need to think of experiences along internalist lines (as involving internal representations) in order to account for illusions such as the phenomenon of apparent motion, and (2) that there is overwhelming evidence of the existence of internal representations in some cases (e.g., to account for the behavior of maze-running rats). This raises a difficult set of topics and requires further exploration elsewhere. A brief response can be made here, however. First, the question at issue here is not whether there are representations in vision, but whether we can assume that to "experience" the environment is to have such representations in one's brain. Second, although it is true that in cases of illusion we do not actively engage with that which we seem to see, it is not the case that in illusions we cease actively to engage with the environment.

[11] As O'Regan reports (based on personal communication with Bach-y-Rita), original experiments involved a stationary camera. It was only when a subject, in frustration, picked up the camera, that he enjoyed "visual experience". Also, of considerable interest: once a

subject has learned to use TVSS, the array of stimulators can be moved from one part of the body to another with no effect on perception.

REFERENCES

Bach-y-Rita, P., C. C. Collins, F. A. Saunders, B. White, and L. Scadden: 1969, 'Vision Substitution by Tactile Image Projection', *Nature* **221**, 963–4.

Bach-y-Rita, P.: 1996, Sustitucion Sensorielle et Qualia, in J. Proust (ed.) *Perception et Intermodalité*, Presses Universitaire de France, Paris, pp. 81–100.

Ballard, D. H.: 1991, 'Animate Vision', *Artificial Intelligence* **48**, 57–86.

Ballard, D. H.: 1996, 'On the Function of Visual Representation', in K. Akins (ed.) *Perception*, Oxford University Press, New York & Oxford, pp. 111–131.

Blackmore, S. J., G. Brelstaff, K. Nelson, and T. Troscianko: 1995, 'Is the Richness of Our Visual World an Illusion? Transsaccadic Memory for Complex Scenes', *Perception* **24**, 1075–81.

Brooks, R. A.: 1991, 'Intelligence Without Reason', *Proceedings of the 1991 International Joint Conference on Artificial Intelligence*, 569–595.

Burge, T.: 1979, 'Individualism and the Mental', *Midwest Studies in Philosophy* **4**, 73–122.

Clark, A.: 1997, *Being There: Putting Brain, Body and World Together Again*, MIT Press, Cambridge, MA.

Clark, A. and D. Chalmers: 1998, 'The Extended Mind', *Analysis* **58**(1), 7–19.

Churchland, P. S., V. S. Ramachandran, and T. J. Sejnowski: 1994, 'A Critique of Pure Vision', in C. Koch and J. L. Davis (eds), *Large-Scale Neuronal Theories of the Brain*, The MIT Press, Cambridge, MA.

Churchland, P. S. and V. S. Ramachandran: 1996, 'Filling in: Why Dennett is Wrong', in K. Akins (ed.), *Perception: Vancouver Studies in Cognitive Science*, Vol. 2, Oxford University Press, New York & Oxford, pp. 132–57.

Crick, F.: 1996, 'Visual Perception: Rivalry and Consciousness', *Nature* **379**, 485–6.

Crick F. and C. Koch: 1990, 'Towards a Neurobiological Theory of Consciousness', *Sem Neurosci* **2**, 263–75.

Crick, F. and C. Koch: 1995, 'Are We Aware of Neural Activity in Primary Visual Cortex?', *Nature* **375**, 121–3.

Crick, F. and C. Koch: 1998, 'Consciousness and Neuroscience', *Cerebral Cortex* **8**, 97–107.

Dennett, D. C.: 1991, *Consciousness Explained*, Little Brown, Boston.

Dennett, D. C.: 1992, ' "Filling-In" Versus Finding Out: A Ubiquitous Confusion in Cognitive Science', in H. L. Pick, Jr., P. van den Broek, and D. C. Knill (eds), *Cognition: Conceptual and Methodological Issues*, American Psychological Association, Washington D.C.

Dennett, D. C.: 1998, 'No Bridge over the Stream of Consciousness', *Behavioral and Brain Sciences* **21**, 753–4.

Feldman, J. A.: 1985, 'Four Frames Suffice: A Provisional Model of Vision and Space', *Behavioral and Brain Sciences* **8**, 265–89.

Gregory, R. L. (1966/1978) *Eye and Brain: The Psychology of Seeing*, 3rd edn, McGraw-Hill, New York.

Grimes, J.: 1996, 'On the Failure to Detect Changes in Scenes across Saccades', in K. Akins (ed.), *Perception: Vancouver Studies in Cognitive Science*, Vol. 2, Oxford University Press, pp. 89–110.

Haugeland, J.: 1995, 'Mind Embodied and Embedded', in J. Haugeland (ed.), *Having Thought: Essays in the Metaphysics of Mind*, Harvard, Cambridge, MA, 1998.
Hurley, S. L.: 1998, 'Vehicles, Contents, Conceptual Structure, and Externalism', *Analysis* **58**(1), 1–6.
Hurley, S. L.: 1999, *Consciousness in Action*, Harvard, Cambridge, MA, 1998.
Intraub, H.: 1997, 'The Representation of Visual Scenes', *Trends in Cognitive Sciences* **1**(6), 217–21.
Levine, J.: 1983, 'Materialism and Qualia: The Explanatory Gap', *Pacific Philosophical Quarterly* **64**, 354–61.
Marr, D.: 1982, *Vision*, W.H. Freeman and Sons, New York.
Milner, D. A. and M. A. Goodale: 1995, *The Visual Brain in Action*, Oxford University Press, Oxford.
Noë, A. and J. K. O'Regan: 2000, 'Perception, Attention and the Grand Illusion', *Psyche*.
Noë, A., L. Pessoa, and E. Thompson: 2000, 'Beyond the Grand Illusion: What Change Blindness Really Teaches Us about Vision', *Visual Cognition* **7**, 93–106.
O'Regan, J. K. and A. Noë: 2001, 'A Sensorimotor Account of Vision and Visual Consciousness', *Behavioral and Brain Sciences* **24** (in press).
O'Regan, J. K.: 1992, 'Solving the "Real" Mysteries of Visual Perception: The World as an Outside Memory', *Canadian Journal of Psychology* **46**, 461–88.
O'Regan, J. K., R. A. Rensink, and J. J. Clark: 1997, 'Picture Changes During Blinks: Not Seeing Where You Look and Seeing Where You Don't Look', *Invest. Ophthalmol. Vis. Sci.* **38**, S707.
O'Regan, J. K., R. A. Rensink, and J. J. Clark: 1999, 'Change-Blindness as a Result of "Mudsplashes" ', *Nature* **398**, 34.
Pessoa, L., E. Thompson, and A. Noë: 1998a, 'Finding Out about Filling-In: A Guide to Perceptual Completion for Visual Science and the Philosophy of Perception', *Behavioral and Brain Sciences* **21**(6), 723–48.
Pessoa, L., E. Thompson, and A. Noë: 1998b, 'Filling-In is for Finding Out (Authors' Response to Commentary)', *Behavioral and Brain Sciences* **21**(6), 781–802.
Poggio, T., V. Torre, and C. Koch: 1985, 'Computational Vision and Regularization Theory', *Nature* **317**, 314–9.
Putnam, H.: 1975a, 'The Meaning of "Meaning", in *Mind, Language and Reality: Collected Papers*, Vol. 2, Cambridge University Press, Cambridge.
Putnam, H.: 1975b, 'Philosophy and our Mental Life', in *Mind, Language and Reality: Collected Papers*, Vol. 2, Cambridge University Press, Cambridge.
Ramachandran, V. S.: 1992, 'Blind Spots', *Scientific American* **266**, 86–91.
Ratliff, F. and L. Sirovich: 1978, 'Equivalence Classes of Visual Stimuli', *Vision Research* **18**, 845–51.
Rensink, R. A.: 2000, 'The Dynamic Representation of Scenes', *Visual Cognition* **7**, 17–42.
Rensink, R. A., K. O'Regan, and J. Clark: 1997, 'To See or Not to See: The Need for Attention to Perceive Changes in Scenes', *Psychological Science* **8**(5), 368–73.
Simons, D. J.: 2000, 'Current Approaches to Change Blindness', *Visual Cognition* **7**, 17–42.
Simons, D. J. and Chakris, C. F.: in press, 'Gorillas in Our Midst: Sustained Inattentional Blindness for Dynamic Events', *Perception*.
Simons, D. J. and D. T. Levin: 1997, 'Change Blindness', *Trends in Cognitive Sciences* **1**(7), 261–7.
Simons, D. J. and D. T. Levin: 1998, 'Failure to Detect Changes to People during Real-World Interaction', *Psychonomic Bulletin and Review* **4**(5), 644–9.

Teller, D. Y. and E. N. Pugh, Jr.: 1983, 'Linking Propositions in Color Vision', in J. D. Mollon and L. T. Sharpe (eds), *Colour Vision: Physiology and Psychophysics*, Academic Press, London.

Thompson, E., A. Noë, and L. Pessoa: 1999, 'Perceptual Completion: A Case Study in Phenomenology and Cognitive Science', in J. Petitot, J.-M. Roy, B. Pachoud and F. J. Varela (eds), *Naturalizing Phenomenology: Issues in Contemporary Phenomenology and Cognitive Science*, Stanford University Press, Stanford, CA.

Trehub, A.: 1991, *The Cognitive Brain*, The MIT Press, Cambridge, MA.

Tsotos, J. K.: 1987, in S. Shapiro (ed.) *Encyclopedia of Artificial Intelligence*, John Wiley, New York.

Varela, F. J., E. Thompson, and E. Rosch: 1991, *The Embodied Mind: Cognitive Science and Human Experience*, The MIT Press, Cambridge, MA.

Van Gulick, R.: 1998, 'Analytical Isomorphism and Marilyn Monroe', *Behavioral and Brain Sciences* **21**(6), 776–7.

Wilson, R.: 1995, *Cartesian Psychology and Physical Minds*, Cambridge University Press, Cambridge.

Wittgenstein, L.: 1953/1958, *Philosophical Investigations*, 3rd edn, Blackwell, Oxford.

Wittgenstein, L.: 1958, *The Blue and Brown Books*, Blackwell, Oxford.

Department of Philosophy
University of California
Santa Cruz, CA 95064
U.S.A.