The Human Nature Review
ISSN 1476-1084
URL of this document
http://human-nature.com/nibbs/03/noe.html

Human Nature Review **3** (2003) 176-182

Book Review

Is The Visual World a Grand Illusion? Edited by Alva Noë. Imprint Academic, 2002 Paperback, 202 pages. ISBN 0907845231.

Reviewed by Paul Coates

This excellent and absorbing collection – a special issue of the *Journal of Consciousness Studies*, edited by Alva Noë, now published in separate book form – comprises fourteen papers dealing with responses to certain very striking recent discoveries about the limitations of visual attention.

A cluster of experiments on "Change Blindness", "Inattentional Blindness" and associated phenomena appear to demonstrate extremely counter intuitive results. According to one plausible characterisation, these results show that we consciously take in far less of the visual world than it seems we are aware of. It is worth briefly summarising the results of two recent sets of experiments, in order to give a flavour of this work. In 'Gorillas in our Midst' (Simons, D. and Chabris, C., *Perception*, 1999, **28**), subjects were asked to perform a task that involved watching a video of a casual basketball game that lasts for about a minute. The task involves counting the number of consecutive passes between members of the one of the teams. While the basketball is being thrown from player to player, something unexpected takes place: a person dressed in a black gorilla-suit walks through the play, stops briefly in the centre of the picture, thumps his chest, and then walks off. Although most subjects correctly record the number of passes made by the team, at least half of these subjects fail to notice the gorilla-suited interloper, who is visible for about nine seconds. When shown the video sequence a second time they are amazed to observe what they had previously overlooked.

In Change Blindness experiments, a picture is displayed of a couple in the foreground with a number of buildings in the background. A brief flicker interrupts the display, which then returns with one substantial change – a large building in the background is no longer visible. This display remains visible for a second or so, and then again the flicker occurs and the switch is made, returning to the original picture. It may take many such alternations between images before the subject finally notices the large scene change not previously attended to; again, this produces surprise on the part of the subjects, when they realise they have overlooked an apparently obvious change. What is agreed by all theorists is that the phenomena of Change and Inattentional Blindness, etc, are well established. These experiments provide strong evidence about the limits of attention. What is less

clear is exactly what they show about the nature of human visual experience.

In his important contribution 'Is Seeing All It Seems?', Andy Clark provides a very useful categorisation of possible responses to the experimental findings. At one extreme is the response that our thinking about perceptual experience, especially about vision, involves what has become known as "The Grand Illusion". Just what the supposed illusion actually *is* needs to be formulated with care. Properly speaking, the illusion relates to what the average layperson assumes about their visual awareness of the objects they see in front of them. The suggestion is that we think we experience much more detail and richness than in fact is present in our experience of the world. We think we notice much more than we really do. The error is, supposedly, about the extent of the surrounding scene that experience enables us to grasp. Clark is careful to distinguish such ideas about the complexity in what experience *presents* to us, from our *reflective* thinking about the nature of experiential states. He thus avoids the trap several other commentators at times fall into, of identifying the Grand Illusion with the view that people mistakenly take their visual *representations* of the world to be rich and detailed. For, as Jonathan Cohen points out in his paper, 'The Grand Grand Illusion Illusion', the notion of a representation belongs to the reflective theorising of Cognitive Scientists and philosophers; any illusion that the average layperson may be under could only concern the contents of their experience, not the ontological status of inner states of mind.

Cohen's own response to the experimental work on attention is deflationary; on his view, the recent results do not really reveal much we didn't already know about limitations on our cognitive and conceptual capacities. In contrast, Daniel Levin argues in his paper, 'Change Blindness: Blindness as Visual Metacognition', that the results are important and novel. Important, because of the implications for visual testimony in criminal cases where eyewitness evidence is pivotal, and also in accident litigation, where unwarranted assumptions may be made about, for example, the visual salience of objects in a driver's field of view. The work is genuinely novel, because it uncovers systematic differences in normal attentional responses to different types of situation, responses which run counter to our traditional assumptions. Perceptual experience, at least at the attentional level, is not what we take it to be. We do not experience what we seem to experience. Put like this, the claim sounds paradoxical. It makes sense to suppose we are mistaken about the theoretical underpinnings of perception and cognition. But how can there be room for mistakes about the nature of our own conscious experience? Any serious attempt to resolve this problem has to appeal to different aspects of conscious experience. One plausible way of doing so involves distinguishing between conceptual and non-conceptual aspects – a distinction motivated on the independent grounds that something must account for the subjective difference between merely thinking about an object in front of oneself, and actually seeing it.

Versions of the Grand Illusion are appealed to in their papers by Daniel Dennett, Susan Blackmore and Bruce Bridgeman, the latter also stressing the top-down processes that affect experience; in a situation where the available input is ambiguous or simply insufficient, the brain will "fill in" so as to construct a plausible representation of the surrounding scene, acting upon assumptions about what the world normally is like. Filling in of certain kinds is also appealed to in papers by Frank Durgin and Charles Siewert. A paper by Eric Schwitzgebel argues that in the related case of visual imagery, there is evidence that knowledge of our own experience is poor.

One of the most interesting lines of response is the sensorimotor theory developed by Kevin O'Regan and Alva Noë in a number of important papers, including their 'A Sensorimotor Account of Vision and Visual Consciousness' (*Behavioural and Brain Sciences*, 2001, **24**). In

his contribution to this volume Noë admirably summarises the experimental work on Change and Inattentional Blindness, and also shows why there are reasons for dissatisfaction with the Grand Illusion response, sketching out his version of the skills theory. In some sense we do seem to be aware of features of the scene that we are not directly attending to. As Nicholas Humphreys has recently emphasised in his *How To Solve The Mind-Body Problem* (2000), if I briefly open my eyes, say for 150 msecs, I get a sense of a whole field of colour and shapes in front of me, even if I cannot immediately discern any particular objects. In Noë's words, 'We are perceptually aware, sometimes, of unattended detail.' This phenomenal fact, the sense of the presence of unattended features in my experience, needs to be accounted for; it won't do to explain it all away as simply arising from a set of confabulated beliefs. For one thing, there is a phenomenal difference between false beliefs and hallucinations (pace Dennett). For Noë, perceptual presence arises in two connected ways. Firstly, from our possession of sensorimotor skills: we have practical knowledge of the way in which the sensory input at our sense organs will vary consequent upon motor activity, knowledge we utilise to guide our actions. Secondly, the information we require to increase the extent of knowledge at any moment about a given feature of the surroundings is already out there, in the environment. We do not have to store huge amounts of data about the whole scene present to us. All we need to do is probe the world to find out more about specific objects in the scene, by changing the direction of our gaze, or moving closer, and so on. According to Noë, '…our sense of the perceptual presence of the detailed world…consists in our access now to all of the detail, and to our knowledge that we have this access.'

Noë illustrates this with examples of amodal perception. If I hold a bottle, I only have physical contact with it at various separate points. However, as Noë points out, I do not merely *think* that the rest of the bottle is present in the way that I think that there is a room next door. I have a sense of the perceptual presence of the whole bottle, which is uninferred. I do not build up an inner representation of it; it is phenomenologically present in my experience by being available to me in the world – I have practical knowledge of the way that the exploratory activity involved in moving my hands down the bottle would give rise to certain patterns of sensory stimulation. Similarly, if I see a cat sitting still behind a picket fence, I have a sense of the presence of the whole cat. The sense of its presence as a unified object 'consists in my expectation that by movements of my body I can produce the right sort of new cat stimulation', and can thereby make out parts of it that I did not see before.

Noë's account provides a convincing analysis of one aspect of the phenomenology of these amodal cases, and is applicable in general to that kind of seeing labelled by Wittgenstein as 'seeing as'. How we organise our experiences is to be cashed out in terms of potential action. Yet it might be objected that the account fails to capture another, different sense of phenomenal presence, the sense in which non-occluded features of the scene are actually in experience, despite being not attended to: for example, when I am aware of the other parts of a Persian carpet in the background while my attention is directed onto the medallion shape at its centre. For Noë, the reply is that the way a feature occurs in experience is always dispositional. Features are present, because they can be accessed in the manner the sensorimotor account spells out.

But 'access' is ambiguous. In one sense, I have access to a feature of the environment when I am having an experience whose sensory content *actually* involves that feature. In another less immediate sense, I have access to information if I am able to get hold of it – I have the information *potentially*, and do not need to be currently in possession of it. The unattended aspects of the carpet, it is arguable, still give

rise to low-level sensory experience, even if there is a lack of detail. In contrast, the unity of the cat features in my awareness at a more conceptual level. To find out more about the cat I have to alter the input to my senses. To find out more about the carpet I have to attend differently to what is already in my experience. The objection is therefore that advocates of the sensorimotor approach conflate two different interpretations of 'access':

> 1) The access a subject currently enjoys, by virtue of what is actually present at a low level in sensory non-conceptual experience, and hence is immediately available for further processing by such activities as change of attention, allowing concepts to be applied to specific aspects of the scene.
> 2) The access a subject has in a dispositional sense, in virtue of being able to generate entirely new sensory input by moving the sense organ relative to the physical object, so that new experiences will result.

Several important issues are raised in a very interesting paper, 'Two Dogmas of Consciousness', by Mark Rowlands, some of which bear upon this problem. Rowlands's paper is marked by a careful consideration of the different aspects of consciousness, and he is one of the few contributors to distinguish clearly between conceptual and non-conceptual states of mind. In the first part of his paper Rowlands argues against what he terms 'objectualism', the view that 'what it is like to have or undergo an experience is an introspectible feature of that experience' – that "what it is like" can be an object of consciousness. Part of the argument here turns on his acceptance of materialism. Nonetheless, as he acknowledges towards the end of his paper, *contra* O'Regan and Noë, there remains the "Hard Problem" of consciousness, and one that on his account there may be little prospect of resolving.

In the second half of his paper Rowlands puts forward an externalist account of experiences, carefully distinguishing such an account from (merely) representationalist accounts of experience. The actual physical objects in the environment that we see make up the contents of our experience. One great advantage of this view is that it gives a nice account of the transparency of experience. The very book I see in front of me is something I am directly aware of, because my experience of it is unmediated. For Rowlands, 'My experience… does not stop short of the public object itself.' Unfortunately, this view is also problematic. It makes sense to claim, as Wittgenstein does, that my thought about my brother in America does not stop short of him. But that is in part because there need be no sensory non-conceptual aspect to my thought. Thoughts are indeed transparent, because they involve no essential phenomenology. It is unclear, however, how this model applies to perception. The reason is that perception involves not only a relation between the subject and object perceived, but also a qualitative subjective awareness of the non-conceptual aspect of experience: sensible properties of the physical object are present in experience in a manner very different from thought.

As Wilfrid Sellars noted in early papers such as 'Phenomenalism' (*Science, Perception and Reality*, 1963), the relation of direct awareness that Rowlands in effect appeals to is the same relation that was traditionally claimed to hold between subjects and their sense-data. The claim used to be that in perception there exists a real relation between two existing entities, the subject, and an existing sense-datum (as opposed to there being a quasi-relation to an intentional, inexistent object). Sense-data theorists were never able to come up with an adequate answer to Ryle's charge: that there is no coherent account of what the relation of direct awareness is. But today's direct realists, such as Rowlands, do not give up on the act-object model of perceiving. Instead they change the object from a private sense-datum to a public

physical item. The relation of awareness is supposed to hold directly between the subject and an existing physical object or event. But can they do any better in explaining what the supposed real relation between act and object amounts to, if they try to avoid countenancing a causal connection between the physical object seen and some inner sensory experience? It is unclear whether the externalist arguments that apply to thinking about individual objects can be simply transposed to perceiving, where the qualitative nature of experience has also to be taken into account.

To his credit, Rowlands does provide an interesting and original attempt to answer this problem. On his version of perceptual externalism – an extended mind account - the objects we see are constituents of perceptual experience because experience extends out beyond the skin of the subject. But problems remain. How are we to understand the basis of the unity involved in the experience - why should the subvening elements involve just the subject and the object in perceptual situations, but not also in other parallel cases, e.g. in hallucinatory cases where there is a non-standard cause of a matching experience? I suggested in 'Perception and Metaphysical Scepticism' (*Supplementary Proceedings of the Aristotelian Society*, 1998), that there is no principled way of answering such questions. Other more traditional problems also threaten to cause difficulties for the view – for when I look at a distant stellar event, it seems we must say that my present experience has as a constituent some event that occurred before I was born. The problem is not the *epistemological* one of justifying our knowledge of what was once in existence, but the *metaphysical* one of explaining how present awareness can encompass events that took place in the distant past. On the face of things, the claim is contradictory.

Andy Clark shares with Rowlands a scepticism about the swift means by which Noë attempts to dispel the hard problem of consciousness, and makes some excellent points about the virtues and drawbacks of the skills theory account. As Clark rightly argues, the hard problem arises for all types of phenomenal experience. It is not clear how Noë's approach could be adapted to deal with our experience of pain, for example; and although Noë's approach certainly sheds light on the differences between experiences, it is still not clear why there should be anything like experience at all for complex systems such as animals, and why animals intuitively differ from simple ping-pong playing robots. He also makes some acute points in arguing that O'Regan and Noë's sensorimotor account ties 'conscious visual experience too closely to the precise details of the low-level sensorimotor routines' by means of which the engagement with the richness of the surrounding scene proceeds. Two subjects whose eyes saccade at slightly different speeds need not have any different visual experiences.

Clark goes on to outline a bold and complex account of his own, combining some of the insights of O'Regan and Noë with the important work of Milner and Goodale on the dual systems hypothesis of visual processing (*The Visual Brain in Action*, 1995). In part, the explanation of perceptually guided action appeals to fairly immediate on-line responses to the environment. For example, responses in sporting activities often require rapidly processed information from the environment for the immediate fine-grained control of action. Such processing may be attributed to the dorsal system. But there is also a role for higher-level storage and 'reason-and-memory-based systems' in selecting the general *types* of actions a subject would perform in a given context. This level, associated with activity of the ventral system, involves conceptual thought and planning, and it is at this level that conscious experience seems to arise. Sameness of experience does not consist in lower level implicit knowledge of sensorimotor contingencies; rather, 'What matters for visual consciousness…. are whatever (perhaps quite high-level) aspects of those sensorimotor contingencies prove most useful for rea-

son, recognition and planning'. Clark's synthesis of the dual systems approach of Milner and Goodale, on the one hand, with the appeal that Noë and O'Regan make to implicit knowledge of the relation of actions to resulting sensory input, offers an exciting prospect for making sense of a range of perceptual phenomena. Nevertheless, on Clark's account there is still the problem of connecting the encoding and storage in the activation of memory structures involved at the conceptual level, with the experience of environmental features. The higher level structures would seem to necessitate some kind of representations, although, as Clark points out, these need not necessarily be of a propositional form. So on Clark's model, the question posed earlier - the question of how we are to make sense of "having access to the information in the environment" - becomes even more pressing.

Two excellent papers, one by Arien Mack, the other by Temre Davies, Donald Hoffman and Antonio Rodriguez, argue in different ways for an alternative approach, one that acknowledges the role of general low-level inner representations of visual input at a pre-attentive stage. Davies et al make the important point that representations need not faithfully reflect all aspects of the surrounding scene. All that is necessary is that a construction of the scene provides a useful guide to future action. The key concept here is survival. If I can act so as to benefit my survival, I may not require full and accurate knowledge of the world. Thus many colour-blind subjects do not realise their condition until it is revealed by careful testing. A doctor of my acquaintance did not realise that he was colour-blind until he attempted to fail aircraft pilots in the RAF on colour vision tests. An angry squadron leader eventually made it clear to him whose eyesight was not normal. Mack reminds us that in many of the experiments on inattentional blindness attention is highly focussed, and background phenomena are ignored. But in other, more normal, contexts, subjects are able very rapidly to ascertain the gist or 'general sense of a complete scene'. We may speculate that this objectively verifiable ability is reflected by the subjective awareness, noted earlier, of the way that a brief glimpse may suffice for an impression of a rich and extensive field of colour.

The appeal to a pre-attentive stage of vision is a central feature in an important theory of attention recently developed by Ronald Rensink, in papers such as 'The Dynamic Representation of Scenes' (*Visual Cognition*, 200, 7 (1/2/3)), and whose work is cited by several contributors. Rensink notes evidence for the formation of a 'low level map-like representation' at an early stage of visual processing. One role of this level is to guide subsequent changes in fixation and attention. It is only at higher levels of processing, where the attention is brought into play, that a 'coherence field' generates the representation (involving higher-level conceptual categories) of a physical object in the environment. Such work suggests a variant of another of the possible responses listed by Clark to the experimental work on Change Blindness, etc., the response dubbed as 'fleeting awareness', or 'inattentional amnesia'. The idea in the more simple-minded formulations of this response is that a subject does indeed take in and process information relating to most of the scene in front of him, but then rapidly forgets it. As Clark notes, this diagnosis results in a rather implausible explanation of the gorilla case. How could observers of the basketball game forget such an extraordinary event as a gorilla walking through the field of play, an event lasting for up to nine seconds? But Rensink's proposal is subtler. The point is that although the sensory information from the whole scene is initially processed to a rudimentary level, little of it is categorised at a higher conceptual level, hence little of it is available for storage in the memory system. Detailed information about specific events in the scene is not forgotten, because most of it is never conceptualised by the subject in the first place. Representations at the low-level are categorised only

by their spatial features and location, and are extremely volatile, being continually replaced without being stored. Nevertheless, such low-level representations still play a crucial explanatory role in accounting for changes in attention, in priming, and in our knowledge of the overall layout, and of the gist of scenes. As suggested above, these representations may contribute to the subjective phenomenology of vision in a way that avoids a problematic appeal to the direct awareness of physical objects. But if there is such a connection with conscious experience, it is unlikely to be straightforward. An interesting paper by Dana Ballard, which concentrates on the different time-scales involved in visual processing, draws attention to the fact that 'there is overwhelming evidence that the interactions between brain and visual world are extremely fragmented.' We can absorb the information from events in the environment for certain purposes rather more quickly than seems to be required for full conceptual awareness of the scene. Just how the various low-level and higher-level visual routines fit together to produce the apparent seamlessness of conscious experience is still little understood.

Overall, this is a most impressive selection of essays on a central topic in cognitive science. Every essay has valuable points to make. Collectively, they show the importance of arriving at a clear conceptual framework for integrating all the different findings about perception. They demonstrate, above all, that in addition to the so-called "easy" and "hard" problems of consciousness there is a crucial third problem – the "conceptual problem" - which consists in clarifying just what consciousness experience itself involves. These essays indicate how difficult, yet how important this third problem is. Between them, they offer a number of interesting suggestions for getting to grips with it. This book is essential reading for any theorist interested in perception.

Paul Coates, Ph.D., Research Leader in Humanities, Department of Philosophy, University of Hertfordshire, United Kingdom. Email: P.Coates@herts.ac.uk.