CHAPTER FIFTEEN

Is the Aim of Perception to Provide Accurate Representations?

Kirk Ludwig

> Ears in the turrets hear
> Hands grumble on the door
> Eyes in the gables see
> The fingers at the locks.
> (Dylan Thomas)

I take the title question to be the question whether the instrumental or biological function of perceptual systems is to provide us with perceptual experiences that are by and large accurate representations of our environments. I will argue that the answer to this question is "yes."

The assumption that perception yields experiences that represent the world around us by and large accurately is deeply embedded in the tradition in the philosophy and psychology of perception. This is the representational theory of perception. On this view, when we perceive objects, we have perceptual experiences, which represent our environments, differently in different modes, but typically through a field medium, as in the paradigmatic case of visual experience, which represents a spatial field filled in in various ways. We form perceptual beliefs on this basis, and our beliefs are accurate, since they in effect abstract from the richer representational medium of experience, just insofar as our experiences are accurate.

This traditional view has been challenged on the basis of recent experimental work in psychology on "change blindness"[1] and "inattentional blindness"[2] among other phenomena, and the accumulating body of knowledge about neurophysiology of vision, which is taken to suggest at the least that our ordinary views about the degree of accuracy and completeness of our visual representations of the world before our eyes are vastly exaggerated, and, perhaps, indeed, that some even more radical overhaul of our traditional picture of perception is required. It has been suggested, in particular, that

> Research in this area calls into question whether we really enjoy perceptual experiences which represent the environment in rich detail. If we do not enjoy experiences of this sort, then we need to rethink the idea that perception is a process aiming at the production of such experiences. (Noë, 2002b, preface)

In recent discussions this has come to be called "the grand illusion," though it is not clear that all participants in the debate express the same view with this phrase.[3]

There are, in this challenge to the traditional view, two different strains to be distinguished. The first and more radical, hinted at in the above passage, raises the question whether perception involves representations at all. If it does not, then the question of their accuracy, and whether they aim at accuracy, does not even come up. The second admits that perception involves representations, but argues that they are largely non-veridical or largely inaccurate. We will take up both sorts of objection in the following.

I will begin by sketching the view that is taken to be under attack, and then sketch some of the evidence that has been advanced against it. Then I take up the radical view, motivated in part by these empirical findings, that denies that perception involves representations at all. I will then return to the question of whether the traditional view and the view under attack are the same, whether we have been in any interesting sense subject to a grand illusion, and whether the empirical findings undermine the view that perceptual experiences are by and large accurate representations of our environment. Finally, I will address in the light of the discussion whether perception aims at accurate representation.

According to many psychologists and philosophers of perception, and even ordinary people – it is said – visual perceptual experience in particular represents the visual scene in uniform, rich detail. This conception of visual experience has been called "the snapshot conception of experience," according to which "you open your eyes and – presto – you enjoy a richly detailed picture-like experience of the world, one that represents the world in sharp focus, uniform detail, and high resolution from the center out to the periphery" (Noë, 2002a, p. 2). Attention ranges over this present rich detail, which is stored in some sense as an internal representation of the environment. It is because of this richly detailed internal representation of the visual scene that we have the impression of a complete and detailed field of vision and a richly detailed world.

On closer examination, however, the snapshot conception of experience begins to look more like a myth. Phenomenologically, it is clear that the visual field (how things are subjectively presented to us in visual experience) is not uniformly detailed and in sharp focus from the center right out to the periphery. Our representation of the scene before us fades as it were toward the edges of the visual field, with far less detail being represented at the periphery than at the center. This is what we should expect given that the sensory receptors on the retina are sparser toward the edge than in the foveal region. We get much less information[4] about what goes on in the regions around us by way of signals generated by light falling on the portions of the retina which are at the periphery. This is strikingly illustrated in a famous experiment, in which subjects were seated before a computer screen and asked to read a page of text while their eye movements were tracked. They had the impression of a

stable page of intelligible text. But outside the area their eyes were focused on, about 18 characters in width, only junk characters were displayed. That is, the intelligible text was shifted with their eye movements, and everywhere else junk characters were displayed.[5] In addition, it is well known that the eye saccades three to four times a second, and during movement there is no information transmitted to the brain. Yet this seems not to be represented in visual experience. We likewise fail to notice the blind spot in the visual field corresponding to the location of the optic nerve on the retina where there are no rods or cones. The information the brain gets, therefore, is sparse over much of the scene before us, intermittent, and gappy. How could experience be any better? And surely the fact that these things do not come to our attention without investigation shows that we had thought experience provided a richer and more accurate representation of our surroundings that in fact it does.

In addition, the recently investigated phenomena of change blindness[6] and inattentional blindness,[7] among others, has been cited as showing that we are radically mistaken about our experience representing everything that goes on in our field of vision. In experiments on change blindness, many subjects fail to notice quite significant changes in the visual scene they are looking at (for example, in the color of flowers and cars or the structure of a house or position of a rail in the background) if the changes are made during eye saccades or during the simultaneous display of visual distractions.[8] More specifically, many subjects fail to report noticing any changes when asked. Subjects thus appear to be blind to changes that occur during such events. Inattentional blindness is a related phenomenon. Subjects fail to notice what seem to be quite significant features of the visual scene if their attention is directed elsewhere. In one study subjects were asked to concentrate on a cross in the middle of a screen presented briefly, then masked, and presented again. They were asked to estimate which arm of the cross was longer. On the third or fourth trial they were presented with a new stimulus, a colored square or moving bar. When the stimulus was presented close to the fixation point, 75 percent of the subjects failed to report it when asked if they noticed anything different. In another widely discussed experiment, subjects were shown a video of two basketball teams, one in black and one in white, each passing basketballs.[9] They were to count the number of passes for one team. They were asked afterwards if they noticed anything unusual. Forty-five seconds into the video an intruder walks through the players, a woman holding an umbrella or a man in a gorilla suit. In some trials the intruder was semi-transparent, and in some fully opaque. Seventy-three percent of the subjects reported nothing unusual in the semi-transparent trials, and 23 percent in the fully opaque trials. The conclusion we are invited to draw is that subjects may fail to see or perhaps visually represent what they don't pay attention to.

I will return to what these phenomena show about the traditional view and their relevance to our question in a moment. I first consider a radical alternative to the traditional view motivated by them, the sensorimotor view of perception advocated by Noë and O'Regan (Noë and O'Regan, 2001). Officially, this view holds that perception does not involve representations at all, and seeks to exhibit it as a matter of a pattern of engagement with the environment. We will be concerned with how tenable it is, and whether it really represents as radical an alternative as is suggested.

The sensorimotor theory holds that "Visual experience [for example] . . . does not consist in the occurrence of 'qualia' or such like. Rather it is a kind of give-and-take between you and the environment" (Noë and O'Regan, 2001, p. 80). "[P]erceivers have sensations in a particular sense modality, when they *exercise their mastery of the sensorimotor laws* that govern the relation between possible actions and the resulting changes in incoming information in that sense modality" (p. 82). It is "in this sense to be 'attuned' to the ways in which one's movements will affect the character of input" (p. 84). It is a "form of practical knowledge" (p. 84). According to the sensorimotor view of perception, then, visual experience is not to be understood in terms of representational states at all, but rather in terms of patterns of behavior and their connection with sensory influx embodied in "the whole neurally enlivened body" (p. 85). This point of view is expressed clearly in the following passage:

> both the representationalist and sensationalist [about visual experience], make a . . . fundamental error. Each relies on a conception of visual experience according to which experiences are internal items of which we become conscious when we undergo them . . . momentarily occurring, internal states of consciousness. . . . As against this conception, I have proposed that perceptual experiences are not internal, momentarily occurring states of this sort. I advocate that we think of experience rather as a form of active engagement with the environment. Perceptual experience is a form of integration with the environment as governed by patterns of sensorimotor contingency. (Noë, 2002c, p. 74)

Sometimes the thesis is put as if it were actual behavior that was crucial – "it is a give and take between you and the environment" ((Noë and O'Regan, 2001, p. 80); "experience is not something that happens in us but something that we do" (p. 99); "sensation occurs when a person *exercises* mastery of those sensorimotor contingencies" (p. 99; emphasis added). I will call this the activity theory. Sometimes the view is put in a way that suggests it is not actual activity that is required, but a disposition of a certain sort. I will call this the dispositional theory. Saying that visual experience is a form of practical knowledge suggests the dispositional, rather than the activity view. And at one point exercising one's mastery of sensorimotor contingencies is characterized as consisting itself in "our practical understanding that if we *were* to move our eyes or bodies" (p. 84) there would be appropriate resulting changes of such things as "the influx from monochromatically tuned rod photoreceptors taking over as compared to influx from the three different cone types present in central vision" (p. 83). Thus, in visual experience, there is a characteristic change in neural influx when we step toward an object and away from it, when we turn our heads to the right or left, or when we close our eyes, or blink. On this view, the sum of all these sensorimotor contingencies and our attunement to them associated with a particular sensation, e.g., of red or yellow, *is* the visual experience or sensation of red or yellow.

The activity interpretation is untenable, as Noë and O'Regan recognize. One may perceive a surface as red so briefly that one has no time to move. Actually engaging in activity cannot be a requirement on perception. Moreover, it seems clearly to be a category mistake to talk of perceptual experience, when we have in mind

something like my visual or auditory experience at a given moment in time, as "something that is performed – enacted – by a living animal engaged in a pattern of exploration of its world" (Noë, 2001, p. 53), or "a form of active engagement with the environment" or "integration with the environment as governed by patterns of sensorimotor contingency" (Noë and O'Regan, 2001, p. 74).[10] A perceptual experience, in the target sense, is a state, not an event. Perceptual experiences change, of course, but these are changes in states. We might as well speak of being red as something that is performed by an object in its response to the various conditions of changing light that affect what it reflects.

The dispositional interpretation appears to be the one intended by Noë and O'Regan, even though their rhetoric often suggests the activity view. However, the retreat from the activity to the dispositional theory robs us of a reason to deny that perceptual experience involves representational states, and so leaves us without a radical alternative to the traditional view. For the dispositional view takes perceiving to be a matter of being in certain states, albeit dispositional states. We can call these experiences if we like, and provide standards of correctness. If one's practical knowledge of sensorimotor contingencies, summed up in one's expectations, is appropriate for one's environment, then the experience is veridical, and otherwise not. Standards of accuracy could be defined similarly.

Apart from not being an alternative to the view that perception involves representations, the dispositional theory has a number of flaws. A disposition has one or more manifestation properties and corresponding manifestation conditions. The manifestation property is exhibited in the appropriate manifestation condition. Salt is water-soluble. The manifestation condition is being put in water. The manifestation property is dissolving. Is having a perceptual experience being in a dispositional state? This is not tenable for a number of reasons. First, it leaves out the phenomenal character of perceptual experiences, which is an occurrent, rather than dispositional feature of us – a manifest property rather than a dispositional one. Second, it is not clear that this view can accommodate hallucinations which we know to be hallucinations. When we have a hallucination which we know to be one, we do not expect any of the usual changes to sensory input given the various possible movements we can make. We may, indeed, have no expectations, implicit or explicit, about what changes in sensory input will occur given various possible movements. Yet, on the sensorimotor view, the consequence would be that we did not have any experience at all. For having the experience is having the practical knowledge of what changes would occur if we were to move in such and such a way. And in this case, we have no such practical knowledge. Third, if we characterize the relevant dispositions in terms of detailed facts about actual human and animal perceptual systems, then it is doubtful that they could be conceptually necessary for perceptual experience. There is no apparent conceptual barrier to creatures constructed quite differently from us having perceptual experiences like ours (Clark, 2002, p. 193ff.). It looks as if the practical knowledge concerned, to meet this challenge, would have to abstract altogether from physical realization. The relevant dispositions would have to be dispositions involving expectations about changing sensory input characterized in terms of its content. But this presupposes an independent characterization of the content of the experiences. Finally, it is not clear that perceptual experience requires even such

dispositions. Consider Galen Strawson's thought experiment involving the weather watchers (Strawson, 1994, ch. 9). The weather watchers are sessile beings. They are unable to move, but we are invited to conceive of them as having perceptual experiences, visual experiences in particular. They watch the passing show, the clouds moving across the sky, the rain, the wind across the grass, the fall of leaves and snow, and so on. The weather watchers are prima facie possible. If the thought experiment is coherent, then it deals a fatal blow to the view that dispositions of the sort that Noë and O'Regan have in mind are necessary for perceptual experience.[11]

I turn to the various considerations which prompted this radical alternative to the traditional view to ask whether they give us any reason to think that though perceptual experiences are representations, they are so misleading or inaccurate that we must call into question whether they could be treated as having the aim of providing us with accurate representations of our environments.

The question whether perceptual experience involves in some way a grand illusion has received starkly different answers in the literature. One author writes, "If [consciousness] seems to be a continuous stream of rich and detailed experiences, happening one after the other to a conscious person, this is the illusion" (Blackmore, 2002, p. 17). The title of another article is "Our perception of the world has to be an illusion" (Ballard, 2002). Another writes: "Is visual consciousness a Grand Illusion? In one sense, the answer must be 'Of course' " (Durgin, 2002, p. 88). But others write just as confidently, "The simplest and most straightforward answer to this question taken literally is no, since, after all, we do see" (Mack, 2002, p. 103). Or, again, "I conclude that once we take care to formulate accurately what we believe, on a first-person basis, about the richness of our ordinary visual experience, efforts to expose this as erroneous on a grand scale collapse" (Siewert, 2002, p. 140). And another calls the affirmative answer to the question "The Grand Grand Illusion Illusion," and concludes that "while [change blindness] and [inattentional blindness] raise a number of interesting empirical questions, the view that they show up a widespread grand illusion concerning perception is itself something of a grand illusion" (Cohen, 2002, p. 141).

There are two different interpretations of the grand illusion hypothesis in the literature. The first – the "world illusion" interpretation – is that perceptual experience does not represent the environment as being the way it is. The second – the "perception illusion" interpretation – locates the illusion not between perceptual experience and the world but between us and perceptual experience. Maybe our perceptual experiences do correctly represent the world around us, but we misrepresent the extent and nature of that representation. Most discussants have in mind the second, but some seem to have in mind at least in part the first (Ballard, 2002; Bridgeman, 2002; Durgin, 2002).

In response to the world illusion interpretation, we can offer a transcendental argument to show that there must be limits to the degree to which our perceptual experiences fail to represent correctly the nature of our environments. It is a condition on the possibility of discovering how our perceptual systems work and the extent to which they do not represent our environment correctly that we come to know quite a lot about our environments. It is a condition on the possibility of our coming to know quite a lot about our environments that, in many cases that we can identify,

our representations of our environment are veridical. Our perceptual beliefs, being abstractions from the contents of our perceptual representations, are correct only to the extent to which our perceptual experiences are veridical. Our perceptual beliefs in turn form the basis for our empirical investigations into the world, our inductive practices, the formation and testing of empirical theories about physical law and the neurophysiology of perception. Any empirical argument to the effect that our experiences did not represent the world by and large correctly, at least with respect to those features that form the basis for our scientific theorizing, would be self-defeating, because its premises would be established by the very perceptual mechanisms it called into question. If it were sound, then there could be no reason to accept its premises.

We could at most be justified empirically in thinking that in some circumscribed respects our perceptual experiences did not correctly represent the nature of the world around us. For any argument in favor of that would presuppose that in other respects, determined by the standpoint from which the skeptical argument was given, our perceptual experiences were largely veridical. One traditional example of this kind of circumscribed skepticism is skepticism about the reality of color, and other so-called secondary qualities. Recently, very general considerations have been advanced to show that the conditions necessary for the success of such a skeptical argument cannot be met (Stroud, 2000, esp. ch. 7). The difficulty is that to identify a general illusion about color, we must simultaneously be able to attribute color experiences and beliefs to people, and to establish that nothing is colored. But the practices which make sense of attributing color experiences and beliefs to people depend upon identifying what they believe relative to the objective features of objects in their environments to which they generally respond. If we can make sense of attributing color experiences and beliefs to people only if we can find those beliefs and experiences to be generally responsive in the right way to colored objects in the environment, then there would be no way coherently and simultaneously to identify color experiences and beliefs and to deny the reality of color. The line of thought here is connected with the application of the Principle of Charity in interpretation, which enjoins one, as a condition on the possibility of finding another person interpretable as a speaker at all, to find him to have beliefs which are about the conditions in the environment that prompt them (Davidson, 2001; Rawling, 2003). If this line of argument can be sustained, then we would have established the stronger conclusion that we cannot show we are mistaken in there being things falling in fundamental categories we represent. The world illusion interpretation of the grand illusion hypothesis, according to which the world that perceptual experience represents to us is largely illusory, or illusory in certain fundamental respects, would be shown to be fundamentally in error.

I turn now to the perception illusion interpretation of the grand illusion hypothesis, according to which the illusion lies in our misrepresentations not of the world but of the character of our perceptual experiences. The perception illusion interpretation is directly relevant to our overall question only to the extent to which the evidence cited calls into question the general *accuracy* of perceptual representations. Let us take up first the challenges to the snapshot model of visual experience. The falsity of the snapshot model, at least if the representations we are interested in

are those embodied in conscious visual experience, is obvious from a moment's reflection. Visual experience does not represent the world in sharp focus, uniform detail, and high resolution from the center out to the periphery. The detail represented in visual experience is not uniform from the center out to the periphery. Even for objects close to one, outside the center of one's visual field the level of detail falls off quite significantly. Even in the center of the visual field not everything is represented in the same degree of detail. Objects nearer or further off than what we focus on are not in sharp focus. Some, like one's nose, are too close to focus on at all. Some are too far (recall Descartes' example of the square tower that looks round in the distance). Things that are in focus for us at a given distance can be brought into sharper focus, up to a limit, by moving closer to them.

Precisely because it is so obvious that the snapshot model does not correspond to the phenomenology of visual perception, however, it seems doubtful that ordinary people have been suffering under an illusion about this. We certainly *behave* as if we think that our visual experience does not represent the world in "sharp focus, uniform detail and high resolution from the center out to the periphery" of our visual fields. If we have an interest in seeing what something is like, *we turn our heads or eyes toward it*, even if it is already visible, and *focus on it*, and *approach it* if necessary in order *to examine it more closely*. If we suffered from a grand illusion because we embraced the snapshot view of perception, we would expect it to show up in our behavior, but it does not. There is no grand illusion we suffer from to the effect that visual experience conforms to the snapshot view of experience.[12]

Does the fact that the visual field is not uniformly detailed, in high resolution, and in sharp focus, over its entire extent, and through its entire depth, show that visual experience does not accurately represent the environment?

In discussing the accuracy or inaccuracy of a representation it is important to keep in mind both the subject and degree of definiteness of the representation. Consider an analogy with maps. An interstate highway map does not misrepresent by failing to represent state routes. Moreover, it has a certain standard of resolution. It is responsible only for features which rise to a certain level of significance. Not every curve in a highway is represented. Therefore, certain actual differences in what is being represented will not be represented by any feature on the map. In this sense, the map fails to represent something that is present, even though it is part of the job of the map to represent things of that sort. But this does not count as a misrepresentation because it is not the map's job to represent to that degree of definiteness. This is shown in our handling of the information which maps give to us. We use them for purposes that do not require greater resolution than they provide. We do not take a topographical map of the United States that represents differences in elevation in 500-foot increments to tell us whether there are any small hills in Florida, and we do not protest that it is inaccurate or non-veridical when we learn that Florida is not completely flat. The lesson carries over to perceptual experience, which has a field-like character.

It is their designers who decide what maps are to represent and what degree of resolution they are to be held to. What determines what visual experience is supposed to represent and what standard of resolution it is to be held to? One might here appeal to the evolutionary function of features of perceptual experience. Yet we

would be convinced that perceptual experience represents our environment even if evolutionary theory were false. The correctness of our judgments about what experience represents and its accuracy is not hostage to our understanding of its biological function. The content of our perceptual representations is given to us by the experience itself. This is largely autonomous from what we believe, as is shown by the possibility of illusions persisting though we realize that they are illusions. An object may still look as if there is a window in it even though we learn that it is the effect of a *trompe l'oeil* painting.

The relevant degree of resolution of perceptual experience for the purposes of assessing it for accuracy is determined by the uses we make of perceptual experience in standard conditions. This includes what beliefs we form on its basis and what we think we need to do to find out more about the scene before us when we have a certain visual experience. To see this, consider a thought experiment. Suppose there were certain plants whose leaves had markings on them, which we discovered by accident could be used as maps of their root systems. What would show what degree of resolution we took them to have? It would be our use of them after some experience with how well features of the markings corresponded with their root systems. The resolution we took them to have would be shown by the limits on the judgments we formed on their basis, what we did and did not take them to show us about the root systems of the plants whose leaves they were. We have of course a great deal of practical knowledge of how well and to what degree of resolution our perceptual experience represents our environments. At a given distance, we know quite a bit about how much detail is visually represented, and how much more we can expect to uncover through closer examination. In addition we deploy a framework of concepts which tells us that even at the greatest resolution we can achieve in optimum conditions in, say, visual experience, there is much detail that escapes our view. In a field of grass, we can see blades of grass at our feet if we look down, but this detail disappears as we look out across the field. This is not to represent the grass fusing into a textured and then smooth green plane as the field recedes from our position. We understand that if we walk across the field, we will see more detail than we did initially. We do not take our visual experience of the field in the distance to be a misrepresentation because it fails to resolve individual blades of grass. We know that at that distance visual experience does not represent to that degree of resolution. If, however, when we started walking we were to find that what we were looking at was not green but brown, then we would conclude that our visual experience had misrepresented what was there. We also understand that the standards of resolution will be attuned to what we focus on, and what portion of the visual field is concerned. The less sharply detailed and focused regions of the visual field away from its center are not more inaccurate representations of those portions of the scene before us, but rather representations that have a lower degree of resolution.[13] At night, when colors are washed out or absent because of the low level of light, we do not take our visual experience to represent objects as colored in shades of black and white, but to fail to be resolving their colors under the illumination conditions. This attunement of our standards of resolution for experience to our practical knowledge guarantees that by and large our perceptual representations do not represent beyond their capacities for resolving detail.

What about the suggestion that visual experience misrepresents because while signals to the brain are interrupted by saccades, our visual experience does not appear to be intermittent, and we do not represent a hole in the visual field where there is a blind spot?

The fact that the signal to the brain is interrupted during saccades does not show that visual experience of objects in the environment is non-veridical or inaccurate any more than the fact that a film of a current event has a maximum frame rate, and so could not be said to be capturing changes in what is being filmed continuously, shows that it is non-veridical or inaccurate. The visual experience is representing the environment, not the mechanisms which implement it. Like a film, it has a maximum sensitivity to change. Changes that fall below the threshold are not represented. That perceptual experience has a limit to its temporal resolution is a matter of everyday experience. If I snap my fingers while watching, I do not see the movement of thumb and finger, only their starting position and end positions. But this does not mean that we misrepresent what happens when things move faster than we can see. When the movement falls below the resolution of visual perception, we fail to represent it, but this is not to misrepresent it.

In the case of the blind spot, there is no question of a misrepresentation with binocular vision, because for each eye the other covers the portion of the visual field it receives no signals from. In the case of monocular vision, the question whether a misrepresentation is involved depends on whether the visual field is filled in in the portion corresponding to the optic nerve or not. If it is, then it is at least sometimes a misrepresentation; if not, it is the absence of representation.[14] Neither case looks to show something significant about whether visual experience generally provides accurate representations of the visual scene. In binocular vision, there is no representational deficit due to the blind spot. At most, in monocular vision, there is a lack of representation of a small area in the visual scene or sometimes a misrepresentation.

Let me turn to evidence for the grand illusion hypothesis drawn from studies of change and inattentional blindness, both of which, I will suggest, are rather ordinary phenomena, and do little to support either the grand illusion hypothesis or the thesis that perceptual experience is inaccurate.

Inattentional blindness, that is, failure to notice or recall things that one was not paying attention to, though these things did clearly physically affect our sensory organs, both intermodally and intramodally, is familiar from ordinary experience. When we concentrate on a visual task – reading, or writing, or painting a design on a cup, we often fail to notice even quite significant aural events in our environment. Similarly, when listening to the radio, or a conversational partner at a cocktail party, we may miss most of what goes on in front of our eyes. And it is a commonplace that one often fails to notice somatic sensation when engaged in a difficult task or one's attention is directed elsewhere – as pickpockets are well aware – even to the extent of not noticing that one has cut or bruised oneself or any sensations associated with that. Likewise, intramodally, one may, in concentrating on what one person is saying, fail to notice what her companion is saying though it is at the same volume. Or one may in keeping track of a sprinter not notice or be able to recall the color of the jersey of the runner next to her or much else about the visual scene. Change

blindness too is a pervasive feature of everyday life. We often fail to notice all the changes in scenes in front of us even as we look at them. Some movie-goers I know have failed to notice that in Luis Buñuel's film *That Obscure Object of Desire* the female protagonist is played by two different actresses. Many card tricks are based on our failure to be able to recall in detail facts about cards we see. Prepare a deck of cards by placing the seven of diamonds and the eight of hearts on the top of a deck of cards, and the seven of hearts and the eight of diamonds on the bottom of the deck. Shuffle the deck in front of the subject without disturbing the two cards on the top and bottom. Ask the subject to take two cards off the top, look at them so that he can recall them, and then place them anywhere in the middle of the deck. Shuffle the cards a number of times, without disturbing the two bottom cards. Place the deck on the table, tap it twice, and then deal the two bottom cards onto the table face up. The subject of the trick will take the two cards dealt out to be those which he had memorized. In this case, clearly it is not a matter of failing to pay attention to the cards which explains why one fails to see that they are not the cards one initially looked at. In drawing attention to these things I do not mean to disparage systematic study of the phenomenon of inattentional blindness and change blindness, but only to point out that it is systematic study of a phenomenon we are already familiar with. If there were a case to be made for a grand illusion involving inattentional or change blindness, it is a case that could be made independently of psychological studies.

What do these phenomena show, first of all, about the extent to which we are subject to an illusion about the completeness of experience? Second, what do they show about the veridicality or accuracy of perception?

In the case of inattentional blindness, it has been claimed that the evidence shows that "there is no conscious perception at all in the absence of attention and therefore no perceptual object can exist preattentively" (Mack and Rock, 1998, p. 227). If this were true, then I think it would be fair to say that we were subject to a kind of illusion that we were conscious of things in our visual or auditory fields about which later we cannot report in much detail. But is it true? Is paying visual attention to something phenomenally like having tunnel vision? Does the rest of the visual field disappear or shrink, so that, except for what you are paying attention to, phenomenally the scene in front of you is just like the scene behind your head? This is an experiment which one can perform without a laboratory, and for my part I can report that it is just not so. I am paying attention at the moment to the words that are appearing on my computer screen as I type. But I do not experience a sudden shrinking of the visual field even if I would not be able to tell you much about the detail of the visual scene outside the area of my attention.[15] Similarly for the intermodal case. In paying attention to the words, my body does not suddenly go numb, I do not suddenly go deaf, etc. It is quite easy to imagine how one's whole experience would be different if in paying visual attention to something one simply ceased to have somatic or proprioceptive or auditory experience. A restricted ability to report on things one is not paying attention to does not impugn the view that if they affect one's sensory organs in a way that usually leads to some effect on the phenomenal character of one's visual or auditory experience, etc., then they have a similar effect on the phenomenal character of the appropriate portion of the visual or auditory

field even when one is not paying attention. For the ability to recall or report that one was in a certain complex phenomenal state and one's being in that state are not the same thing, and it is no surprise that we are better able to recall matters involving, and experiences of, things we are paying attention to than things we are not paying attention to.

Why would anyone suggest attention was necessary for consciousness? Mack and Rock reach their conclusion by identifying being able to recall or report something later with having at the time been conscious of it: "A perception is . . . conscious if subjects can recall or recognize it" (Mack and Rock, 1998, p. 233). However, while this is plausibly a sufficient condition for having been conscious of it, it is not a necessary condition, at least if we mean conscious or phenomenal experience. One could defend Mack and Rock's conclusion by introducing an operational definition of "conscious," which does not aim to capture the ordinary meaning, and is tailored to their experimental results. But the air of excitement goes out of the announcement when we take "no conscious perception" to be shorthand for "no or limited ability to recall in the experimental conditions."

This point applies equally to change blindness. Change blindness does not directly show that we do not at time t and at time $t + \varepsilon$, after the change, represent correctly features which have changed. What is shown at most is that one may have limited ability to notice a change in the scene, and by extension in the representation.[16] For change in the world is represented in experience by a corresponding change in what the experience represents, and so in the experience itself. If an object is blue at one time, then red, one's experience represents that change if before the change it represented the object as blue and after the change it represented it as red. To notice that one's experience has represented a change requires taking note of a difference between one representation and another. The results of change blindness experiments do not suggest that before and after the change one's experience does not correctly represent. So they do not suggest that one's experience does not represent a change. The experimental results suggest only that we may fail to notice changes in our experience when they occur during saccades, or blinks, or when there are simultaneous distracting events in the visual scene. Given this, it is a mistake to suppose that people thinking that they would notice such changes shows that they are subject to an illusion about the accuracy or veridicality of their experience.[17] Rather, they overestimate the extent to which they are able to attend to changes in their experience, and remember the character of their experiences at later times.

It is easy enough to explain why we take ourselves to be better at noticing changes in the special situations that elicit change blindness. As Jonathan Cohen has noted, "all the inductive evidence available to perceivers supports their belief in their ordinary capacity to notice ambient events" (Cohen, 2002, p. 152). We typically do notice changes in our environments that are important to us. It is natural then that we should be surprised when we fail to notice some changes that in retrospect seem obvious in circumstances that we do not know are rigged to test the limits of our abilities.[18] But as Cohen remarks, this should no more incline us to say we are subject to a grand illusion than the fact that we are surprised that we are mistaken in the Müller-Lyer or Ponzo or Ebbinghaus illusions. The "grand illusion"

> is an instance of a very general phenomenon: ordinary subjects are ignorant about the limitations on their cognitive and perceptual capacities, and when controlled experimental conditions make these limitations apparent, they (and we) learn something new. (Cohen, 2002, p. 155)

Given that perceptual experience does by and large provide accurate representations of the environment, the question whether that is its aim is straightforward. Experience has the instrumental function of providing accurate representations if its doing so helps us achieve our aims. It is clear that knowing about the environment is important to our achieving many of our aims. This requires correct perceptual beliefs about the environment. And since these abstract from our perceptual experiences, this requires accurate perceptual experiences. Accurate perceptual experience therefore helps us achieve our aims. Perception therefore has the instrumental function of providing accurate representations. Any answer to the question of whether the biological function of perceptual experience is to provide accurate representations is more speculative, since it is an empirical question whose confirmation depends upon historical facts we have only indirect evidence about. Yet it seems overwhelmingly plausible that accurate representations of the environment tailored to an organism's needs provides a selectional advantage. Given this, we may safely conclude that it is also a biological function of perceptual experience to provide accurate representations.

Notes

1 See O'Regan et al., 1996, 1999; Simmons, 2000; Simmons and Levine, 1997.
2 Mack and Rock, 1998.
3 The phrase was introduced into the literature in Noë et al., 2000. A recent issue of *The Journal of Consciousness Studies* (Noë, 2002b) has been devoted to it.
4 There is a dangerous ambiguity in "information" which it would be well to note here. In the text, I use "information" in the sense of a physical signal which together with appropriate laws and background conditions enables one who knows the laws, background conditions, and signal, to infer something about its cause. In this sense, rings in the trunk of a tree carry information about its age. This does not mean that they carry information in the sense in which a newspaper does. A newspaper carries information in two senses, in the signal sense, and in the sense that it represents that certain things have occurred, that is, it contains representations that have intentional content and are true or false. Tree rings are not intentional and are not true or false.
5 O'Regan, 1990.
6 See Simmons, 2000 for a recent review.
7 See Mack and Rock, 1998.
8 O'Regan, 1992; O'Regan et al., 1996.
9 Neisser, 1979; Simons and Chabris, 1999.
10 "Experience" has an event as well as a state reading. However, it is the state reading which is at issue in the question whether the sensorimotor view provides an adequate analysis of perceptual experience in the sense in which we speak of my visual or auditory experience at a given moment in time.
11 The same point can be made by the more traditional thought experiments of the brain in a vat, and the disembodied mind all of whose experiences are determined by an evil demon.

12 Consider in this respect Daniel Dennett's example of walking into a room and seeing wallpaper, in his example, of identical photographic portraits of Marilyn Monroe. Dennett says that "you would see in a fraction of a second that there were 'lots and lots of identical, detailed, focused portraits of Marilyn Monroe' ", but that since "your eyes saccade four or five times a second at most, you could foveate only on one or two Marilyns in the time it takes you to jump to the conclusion *and thereupon to see* hundreds of identical Marilyns" (Dennett, 1991, p. 354). Dennett says rightly that we do not represent in detail more than we actually foveate on. But then he goes on to say: "Of course it does not seem that way to you. It seems to you as if you are actually seeing hundreds of identical Marilyns" (Dennett, 1991, p. 355). But this needs to be handled carefully. You do see a wall on which there are hundreds of portraits of Marilyn Monroe which are detailed. And it seems to you as if you do. But does it or would it seem to you that your *visual experience represented all of that detail*? I don't think that anyone would be under the illusion that it did. It is just that we know that wallpaper involves repetition of a pattern, and if we see the pattern, we know we are seeing a wall on which the pattern is repeated in all its detail. There is no illusion, and no surprise, in any of this.

13 Space constraints prevent a detailed discussion of Kathleen Akins's interesting argument that the peripheral thermoreceptor system does not provide veridical representations (Akins, 1996). The argument is based on an observation and an assumption. The observation is that the warm and cold spots that respond to temperature and temperature change are distributed unevenly over the skin, and have both static and dynamic responses that are nonlinear. The assumption is that intensity of felt sensation represents surface skin temperature if anything. It is the assumption that I would question. We treat sensations of heat and cold as providing information about distal objects and objects we are in contact with, not our skins, and, as in the case of visual experience, the relation between the subjective features of experience and the representation of objectively unvarying properties in the environment may be quite complex. The angle subtended by an object on the retina is not directly correlated either with its represented size or shape, which depends in addition on the represented distance and viewing angle. We may look for a similar interplay between what is represented and a variety of different sorts of information, including cross temporal information, in the case of sensations of heat and cold. For example, when we step into a hot bath, we know that the intensity of the sensation of heat will diminish after a moment. But we do not take this to be a representation of the bath water cooling down – we do not suddenly plunge the rest of the body in after the feet have ceased to complain.

14 Dennett claims there is no filling in Dennett, 1991, but see Pessoa et al., 1998 for discussion and some contrary evidence.

15 Fortunately, drivers do not go blind when they are talking on a mobile phone, though they are apt to do very poorly in reporting on the visual scene before them.

16 See Simons et al., 2002 for some recent experimental work that suggests under probing subjects often can recover information about a scene that it seemed initially that they had not taken note of. This suggests that "change blindness" as defined operationally in these experiments does not correspond to failure to be able to recall and report on the change at all, but failure in response to open-ended questions to make comparisons that would have called the change to mind.

17 See Levine, 2002 for studies of the extent to which people overestimate their ability to detect change.

18 In an informal survey I have found that people have difficulty picking out the difference between the pair of photographs reproduced in Blackmore et al. 1995 when viewing them at the same time, though they pick out the difference easily on being told what to look

for. No wonder subjects can fail to notice a change when they are presented one after another with an intervening eye movement.

References

Akins, K. (1996). Of sensory systems and the "aboutness" of mental states. *Journal of Philosophy*, 937, 337–72.

Ballard, D. H. (2002). Our perception of the world has to be an illusion. *Journal of Consciousness Studies*, 9/5–6, 54–71.

Blackmore, S. J. (2002). There is no stream of consciousness. *Journal of Consciousness Studies*, 9/5–6, 17–28.

Blackmore, S. J., Brelstaff, G., Nelson, K., and Troscianko, T. (1995). Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes. *Perception*, 24, 1075–81.

Bridgeman, B. (2002). The grand illusion and petit illusions: interactions of perception and sensory coding. *Journal of Consciousness Studies*, 9/5–6, 29–34.

Clark, A. (2002). Is seeing all it seems? Action, reason and the grand illusion. *Journal of Consciousness Studies*, 9/5–6, 181–202.

Cohen, J. (2002). The grand grand illusion illusion. *Journal of Consciousness Studies*, 9/5–6, 141–57.

Davidson, D. (2001). Radical Interpretation. In D. Davidson, *Inquiries into Truth and Interpretation* (2nd edn.). New York: Clarendon Press.

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

Durgin, F. H. (2002). The Tinkerbell effect: Motion perception and illusion. *Journal of Consciousness Studies*, 9/5–6, 88–101.

Levine, D. T. (2002). Change blindness as visual metacognition. *Journal of Consciousness Studies*, 9/5–6, 111–30.

Mack, A. (2002). Is the visual world a grand illusion? *Journal of Consciousness Studies*, 9/5–6, 102–10.

Mack, A. and Rock, I. (1998). *Inattentional Blindness*. Cambridge, MA: MIT Press.

Neisser, U. (1979). The control of information pickup in selective looking. In A. Pick (ed.), *Perception and its Development*. Hillfield, NJ: Erlbaum.

Noë, A. (2001). Experience and the active mind. *Synthese*, 129/1, 41–60.

— (2002a). Is the visual world a grand illusion? *Journal of Consciousness Studies*, 9/5–6, 1–12.

— (ed.) (2002b). *Is the Visual World a Grand Illusion?* Special issue of *Journal of Consciousness Studies*, 9/5–6. Charlottesville, VA: Imprint Academic.

— (2002c). On what we see. *Pacific Philosophical Quarterly*, 83/1, 57–80.

Noë, A. and O'Regan, J. K. (2001). What it is like to see: A sensorimotor theory of perceptual experience. *Synthese*, 129/1, 79–103.

Noë, A., Thompson, E., and Pessoa, L. (2000). Beyond the grand illusion: What change blindness really teaches us about vision. *Visual Cognition*, 7/1–3, 93–106.

O'Regan, J. K. (1990). Eye movements and reading. In E. Kowler (ed.), *Eye Movements and Their Role in Visual and Cognitive Processes*. Amsterdam: Elsevier.

— (1992). Solving the "real" mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, 3/46, 461–88.

O'Regan, J. K., Rensink, J. A., and Clark, J. J. (1996). "Mud splashes" render picture changes invisible. *Investigative Ophthalmology and Visual Science*, 37, 213.

—, —, and — (1999). Change-blindness as a Result of "Mud splashes." *Nature*, 398, 34.
Pessoa, L., Thompson, E., and Noë, A. (1998). Finding out about filling in: a guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences*, 216, 723–802.
Rawling, P. (2003). Radical interpretation. In K. Ludwig (ed.), *Donald Davidson*. New York: Cambridge University Press.
Siewert, C. (2002). Is visual experience rich or poor? *Journal of Consciousness Studies*, 9/5–6, 131–40.
Simmons, D. J. (2000). Current approaches to change blindness. *Visual Cognition*, 7/1–3, 1–15.
Simons, D. and Chabris, C. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, 28, 1059–74.
Simmons, D. J. and Levine, D. T. (1997). Change blindness. *Trends in Cognitive Sciences*, 1/7, 261–7.
Simons, D., Chabris, C., Schnur, T., and Levine, D. T. (2002). Evidence for preserved representations in change blindness. *Consciousness and Cognition*, 11/1, 78–97.
Strawson, G. (1994). *Mental Reality*. Cambridge, MA: MIT Press.
Stroud, B. (2000). The Quest for Reality: Subjectivism and the Metaphysics of Colour. New York: Oxford University Press.