READINGS IN MIND AND LANGUAGE

1 Understanding Vision: An Interdisciplinary Perspective
Edited by Glyn W. Humphreys
2 Consciousness: Psychological and Philosophical Essays
Edited by Martin Davies and Glyn W. Humphreys

Consciousness

Psychological and Philosophical Essays

Edited by
Martin Davies and
Glyn W. Humphreys

13

A Bat without Qualities?

KATHLEEN A. AKINS

The Bird's Eye View

The other day in a physiology seminar we were discussing the effect of retinal foveation on visual perception. The fovea is a small portion of the retina densely packed with receptor cells – a density that makes possible those visual tasks that require high spatial resolution, the identification of shape and texture, accurate depth perception and so on. The fovea, however, can 'see' only a small part of the entire visual field. So, much like directing a telescope across the night sky, foveated creatures move their eyes – shifting the 'interesting' parts of the scene in and out of the foveal area. This is why we, but not rabbits, move our eyes about.

Enter the eagle – or, rather, birds of prey in general. They too have foveated eyes, but eyes with even better spatial resolution than our own. The African vulture, for example, can discern live prey from dead at an elevation of 3,000–4,000 metres, an elevation at which it is difficult for us even to sight the bird (Duke–Elder, 1958). Eagles, too, have high resolution foveae. Because they dive for the ground at speeds greater than 200 mph, their eyes must be capable of extremely accurate depth perception. Indeed, given the broad range of visual information that an eagle makes use of in its behaviour, the evolutionary 'solution' was the development of *two* circular foveae connected together by a horizontal band of densely packed receptor cells (think here of the shape of a barbell). The horizontal band serves to scan the horizon. The central fovea, like those of most birds, looks to either side, each one (in the left and right eyes) taking in a different part of the world. Finally, the eagle has an extra pair of (temporal) foveae pointing forward, converging on a shared field – a foveal pair much the same as our own except with three times the density of receptor cells (Duke–Elder, 1958). It is this forward-looking foveal region that provides the high spatial resolution. Attending to

the scene below via the temporal fovea, eagles spot their prey and dive at fantastic speeds, pulling up at exactly the right instant.

But therein lies a mystery, I thought, the mystery of the 'eagle's eye' view. Given two foveal areas and a horizontal band, how does an eagle 'attend to' a scene, look at the world? What does that mean and, more interestingly, what would that be like? Here, in my mind's eye, I imagined myself perched high in the top of a dead tree sporting a pair of very peculiar bifocal spectacles. More precisely, I pictured myself in a pair of quadra-focals, with different lenses corresponding to the horizontal band, foveal and peripheral regions of the eagle's eye. I wonder whether it is just like that, I thought, like peering successively through each lens, watching the world move in and out of focus depending upon where I look. First I stare through the horizontal section and scan the horizon for other predators; then I switch to my left central lens and make sure no one is approaching from behind; then I use the high-powered temporal lens to scrutinize the water below for the shadows of some dinner. Is that how the world looks to an eagle?, I wondered. Is that what it is like to have two foveae?

The Problem: Nagel's Claim and its Intuitive Basis

In 'What is it like to be a bat?' (1974), Thomas Nagel made the claim that science would not, and indeed, could not, give us an answer to these kinds of questions. When all of science is done and said – when a completed neuroscience has told us 'everything physical there is to tell' (Jackson, 1982, p. 127) – we will still not understand the experiences of an 'essentially alien' organism. It will not matter that we have in hand the finer and grosser details of neuroanatomy, neurophysiology and hence, the functional characterization of the system at various levels of complexity – nor will the 'completed' set of psychophysics provide us with the essential interpretative tool. For all of neuroscience, something would be missed – what it is like to be a particular creature, what it is like *for* the bat or the eagle.

There are many reasons, I think, both intuitive and theoretical, why Nagel's claims about the limits of scientific explanation have seemed so plausible. Nagel himself, for example, argued for this conclusion by appeal to a theoretic notion, that of a point of view. Phenomenal experience, he said, is necessarily an experience from a particular point of view, hence the facts of experience are essentially subjective in nature. On the other hand, the kinds of phenomena that science seeks to explain are essentially objective, or viewer independent – 'the kind [of facts] that can be observed and understood from many points of view and by individuals with differing perceptual systems' (Nagel, 1974, p. 145). So any attempt to understand the experience of an alien creature by appeal to scientific facts (facts about his behaviour and internal computational/physiological processes) will only serve to distance us from the very property we seek to explain: the subjectivity of phenomenal experience. Or so Nagel argued. Nagel's conclusion was that the only possible

access one could have to the phenomenal experience of another organism is by means of a kind of empathetic projection – by extrapolation from one's case, we can ascribe similar experiences to other subjects. Needless to say, this is a process that will work well enough given a suitably 'like-minded' organism (such as another person) but which will be entirely inadequate for understanding the point of view of more alien creatures. Hence, given only empathetic means, said Nagel, we cannot know the nature of a bat's phenomenal experience.

Nagel's argument, like those of a number of other philosophers (for example, see McGinn, 1983), makes use of a variety of theoretic tenets – about the objectivity of scientific facts, the subjectivity of experience and about the nature of a point of view. In the usual case, such arguments hinge upon a claim that 'you can't get from there to here' – that there is no route from the objective to the subjective, from the non-intentional to the intentional, from the sub-personal to the personal, and so on – even given all of the resources of the natural sciences. These are views that must be addressed, I think, by argument, each in its own right or, better, met by a demonstration that the dichotomy at issue can in fact be bridged by scientific insight. Rather than address here these theoretic concerns, about subjectivity, point of view and so on, I want to look instead at the *intuitive* pull towards Nagel's conclusion – why most of us harbour that nagging suspicion that science must fail, that it cannot tell us what we want to know. This is the intuition that science will necessarily omit the one essential element of phenomenal experience, namely its very 'feel'.

The unfortunate fact of the matter, I think, is that these negative intuitions are well grounded in our everyday experiences. We have all faced the difficulty of trying to communicate the nature of a particular phenomenal experience, good or bad. 'It was awful, absolutely horrible!!' you might recount, speaking of a bad migraine headache – but, apart from a fellow migraine sufferer, no one seems the wiser for your description. Frustratingly, despite the listener's own extensive catalogue of aches and pains, any elaboration on the 'horribleness' of a migraine seems to do little good. 'Yes, it's a bit like that but . . . .' one will hedge, when asked how a migraine compares to an ordinary headache, one caused by tension or by sinus inflammation. Or is it like having a nasty hangover, a bad case of the flu, or like the stabbing pain one feels when the lights are suddenly switched on in a darkened room? 'It's sort of like that, except, only, um . . . well . . . much, much *worse*!' This is what a sufferer will typically reply, unsure, even in his own mind, what to make of such comparisons. (Does a migraine differ from a bad hangover only in intensity or is there in fact a difference in kind? Or does the difference in intensity *constitute* a difference in kind?) Ironically, the best descriptions one can give, the descriptions that elicit the most empathetic sounds and nods, are usually not descriptions of the pain at all, but of the beliefs and desires that go along with the migraine. 'If I knew the migraine wasn't going to end, I'd seriously wonder whether life was worth living' or 'the pain is so intense, you don't even want to roll over, to find a more comfortable position in which to lie' – it is such thoughts that make clear the severity of the experience. Describing the feelings *per se* just does not seem possible. You simply have to have a migraine.

Extend, then, this epistemic difficulty to the phenomenal experience of an alien creature. Suppose that an organism has sense organs of a completely unfamiliar kind and, further, that it processes the information gathered from these strange sense organs in a manner unique to its species (or at least, in a manner unknown to ours). This is an organism that, undoubtedly, will have experiences that we do not: some of its sensations will be nothing like our sensations. So if we think of an organism's phenomenological experience as constituted by the set of all those alien 'qualia', the problem of understanding seems insuperable. Given that we cannot comprehend by description the relatively familiar and circumscribed sensations of the migraine sufferer, what could we possibly know about an alien creature's point of view – about an entirely foreign phenomenological repertoire? If we can comprehend only those sensations that we have experienced, and if our own sensations are very unlike those of the bat, then we will be unable to understand a bat's phenomenology. This is the intuitive conclusion grounded in everyday experience.

The problem about the experience of bats, however, was, as Nagel described it, a problem about scientific description – whether science, not everyday conversation, could buy us any leverage on the bat's point of view. So what does common sense tell us here? The answer, I think, is that our conclusions about the ineffable nature of sensations fit hand in glove with another common feeling about the efficacy of science: to the average person, the suggestion that science might resolve these communicative difficulties seems quite strange, if not downright puzzling. How could science possibly help us in this respect?

Suppose, for example, that I am trying to describe to you a certain kind of feeling, say the pain of my broken toe. I might say something like this:

> Well, at first, when I tripped over the broom handle, there was a sharp, intense pain – a blinding flash of 'white' that occurred behind my eyes. Then the pain evened out to a dull throbbing in the toe – and, later, by that night, it had turned into what I think of as 'pain somewhere'. You know, that's the pain of a deep injury – when the pain is clearly where it should be, in this case, in the *toe*, but it's also nowhere in particular. Your whole body feels, well, dragged out.

If you have actually had a broken toe or another injury of this sort, these sensations might sound quite familiar. You know, for example, exactly what I mean by the phrase 'a blinding pain'. But if you have been fortunate enough to have avoided such traumas, certain parts of the description will seem quite peculiar. (A 'throbbing' pain you can understand, but what is it to have a pain that is 'blinding' or felt 'nowhere in particular'? Surely this is just a figure of speech?) One can, of course, on the basis of the description, obtain

some understanding of the phenomenological properties at issue (after all, if asked about the pain of a broken toe, you could simply paraphrase the above description!). But it does little to help you understand how the pain actually *feels*. That is the part you cannot grasp given the description alone. Imagine, now, that you are given a completed model of human nociception, a model of all the neurophysiological/computational processess that underlie the production of pain, including, of course, the pain of a broken toe. That this model could in any way help seems entirely dubious. Why would you understand the *pain* of a broken toe any better if presented with a corpus of facts about C-fibres and A-fibres, conductance times, cortical and sub-cortical pathways, transmitter release, the function of endogenous opiates and so on? How could these statements about brain function possibly tell you about the *feeling* of a broken toe?

It is this intuitive sense of puzzlement, I think, that lies behind the more theoretical philosophical arguments of Nagel (1974), Block (1978), Jackson (1982), McGinn (1989) and Levine (chapter 6, this volume) – behind philosophical arguments that 'you can't get from here to there', that there is an unbridgeable explanatory gap between the facts of science and those of subjective experience. In this sophisticated guise, the puzzlement is not given a naive dualist expression: most philosophers do not hold that science must fail to explain phenomenological events because those events occur in a 'realm' beyond the physical world. Rather, the materialistic tenets are upheld: descriptions of neurological processes, it is generally agreed, are descriptions of inner sensations *in some sense of the phrase*. Moreover, given that sensations *are* brain processes, most Nagelians admit that science could not be entirely irrelevant to our understanding of an alien creature's experience. Neurophysiology, psychology and psychophysics will illuminate (no doubt) some aspects of an alien point of view. Still – and this is where the intuitive puzzlement resurfaces – no matter how much we come to understand about a brain's representational or computational capacities (the nature of its functional states at various levels of description, plus their structural and relational properties), the qualitative properties of that organism's point of view will still be missing. Again, it is the 'very feel' of the experience that science is said to leave out. But what exactly does this mean? What is given and what is not by science?

Think here of the difference between, say, a pristine page in a child's colouring book, with only the thick black outlines of the picture drawn in, and that same page alive with colour, the trees and flowers and birds given hue according to the whims and palette of a particular individual. In one we have the 'basic outline' of the image, the two-dimensional form; in the other, we have that outline plus the hues of the forms – colours that might have been different had the artist chosen otherwise. Now if we were given only the pristine page, various questions about the scene would remain unanswered. 'But is the sky blue or is it really grey?' 'Is the flower on the left yellow or is it actually white?' Without the completed picture, it is impossible to tell. It is questions analogous to these, then, that are allegedly left

unanswered given only the neurological/computational facts about another organism's brain processes. Even if we knew the basic outline or, in Nagel's terms (1974, p. 179) the 'structural properties' of an alien creature's representational scheme, the very 'colour' of the experiences, the qualia, would still be missing. Like the missing colours of the outlined page, there are any number of ways, consistent with the structural properties of the representations, that those subjective experiences could be. What science can give us, at best, are boundaries on the space of possible qualia, on the pure 'colours' yet to be filled in. In this way, our everyday intuitions cast the problem of consciousness, both in its naive and philosophical forms, as largely a problem about the intrinsic or qualitative nature of sensations, about the 'greens, reds and blues' of phenomenal experience.

The Film

Imagine, then, that I, having dropped in from some future time towards the end of neuroscience, claim to have a film of 'what it is like'. I have, that is, a film of the phenomenology of the bat. While such a suggestion might at first seem unlikely, let me assure you that this film carries the stamp of approval of future science. For what science has found out, in the fullness of time, is that just as some people have suspected (Dawkins, 1986), the bat's sonar echo is used to solve the very same informational problems for which we humans use light. The bat uses the informational properties of sound to construct a representation of objects and their spatial relations. This is why the bat's experience can be presented on film to us, the human observers – why it has, I claim, a strangely 'visual' quality. Needless to say, this film was made in the appropriate Disney style: a 'cinerama' or 'sen-surround' film projected on a curved screen, 180 degrees around the theatre, presented to an audience outfitted in '3-D glasses', for the sake of stereo vision. And, of course, the film is in colour.

What, then, does the bat film look like? First, the plot is simple. It shows, from the bat's auditory viewpoint, a boring sort of chase scene: the bat, flying about, uses sonar signals to catch mealworms that have been thrown into the air by an experimenter. (Bats, of course, are not blind – they see as well as hear. For the purposes of this thought experiment, however, I am considering only their auditory sensations.) This feat is accomplished with a manoeuvre characteristic of the Little Brown bat. First the bat flaps around, emitting his Fm sonar signal (a cry that begins at about 60 khz and sweeps downward, through the intermediate frequencies, to a cry of about 20 khz) and waiting for something edible to appear; then when he sights a mealworm, he flies over and manoeuvres until he can swat the mealworm with his wing; performing a somersault, the bat then secures the prey in his tail pouch; finally, he reaches down to grab it, eating the mealworm from his pouch (figure 3). (Why bother with the pouch? As someone recently pointed out, 'Every good meal deserves to be eaten sitting down.'[1]) This is the basic

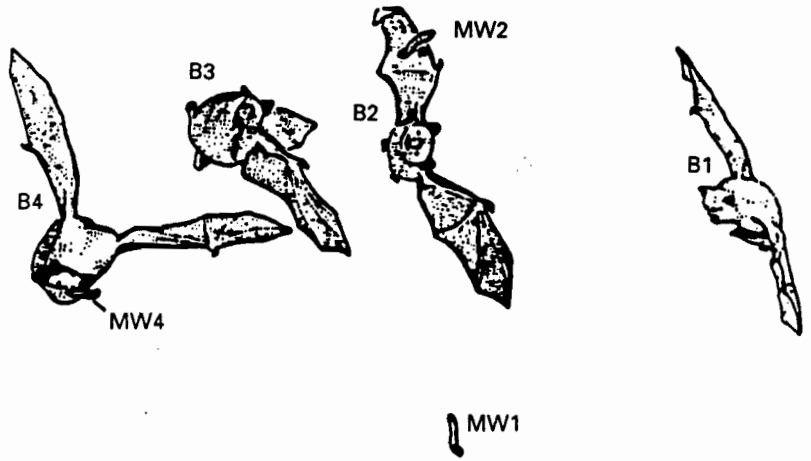

Figure 3 *A filmed sequence of a bat (*Myotis lucifugus*) capturing, by a somersault manoeuvre, a mealworm tossed into the air. Shown are the four sequential positions, beginning with the rightmost figure. In the first frame, the bat (B1) spots the mealworm (MW1), which is still rising from the toss. In the second frame, the bat uses its wing to deflect the worm downward. Next the bat catches the mealworm in a pouch between its tail and two legs. Finally, the bat ducks down to scoop out its meal. (Adapted from Webster and Griffin, 1962.)*

scenario, one that is repeated several times. Now, what the film actually shows to the human observer is a kaleidoscopic display of vibrant colour forms. Swirling and pulsating in three-dimensions, the coloured forms dance across the screen, colliding and dispersing, suddenly appearing or vanishing. That's all. That, I claim, is what it's like. It is not, of course, what we humans would see, if we were acting the part of the bat – if we, with our human visual systems, were trying to catch a mealworm (Nagel 1974). It is not 'visual' in the human sense. On the other hand, this is not a film from our point of view, but from the point of view of a bat.

As you, the reader, will no doubt object, something is clearly wrong with this story. That is, whether or not the film 'accurately depicts' some part of the bat's phenomenology – the sensory 'colours' – watching the swirling display seems to leave out much of what is surely important to the bat's point of view. First, unlike our experiences during a film of a roller-coaster ride or a hang-glider's flight, we do not feel any of the additional 'sympathetic' sensations appropriate to the moment. It does not seem to us that we are making any of the swooping and diving movements that are made by the bat. Nor do we understand the significance of the coloured images. Barring

any sub-titles of the form 'now the somersault begins' or 'now you've got the mealworm in your pouch', you will not know what is happening – what you, as a bat, are doing. When the bright red image swirls across your left 'auditory' field, is something (the mealworm? a background object?) moving past you or are you moving relative to it (maybe this is a somersault?)? Then again, is anything even moving at all? Can you infer that the movement of the colours stands for movement in the world? Probably not. And what does the three-dimensional nature of the film buy you? What does it mean when one coloured patch appears behind or in front of another? Is this a spatial relation or . . . .? All in all, the coloured images hold little insight for the human observer.

As a first pass at explaining what is wrong with this story – why a cineramic film could not tell us what we want to know about the bat – note that, while not particularly helpful in this instance, such 'sen-surround' films are extremely useful in understanding the human point of view. When we watch a film of, say, the hang-glider's flight, the pictures go proxy for the real world. The brain interprets the intensity, frequency and spatial cues of the film in much the same way as it would interpret these same properties of light, reflected by real objects in the three-dimensional world. Hence, we really do see (more or less) what is seen during a hang-glider's flight. Indeed, because the visual system informs both the vestibular and the sympathetic nervous systems, we even feel the non-visual sensations – the terror before the leap, the drop in the stomach that follows. Through watching the film, seeing from this novel perspective the world rush by and feeling the sympathetic sensations of movement, a good deal about the experience of hang-gliding is communicated. In other words, we can simulate another person's point of view just because (a) we share a similar visual system, and (b) we can artificially create the hang-glider's visual input.

Similarly, when we watch the film of the 'bat experience', we use the spectral cues in ways typical of human vision (what other choice could there be?). But what exactly does that mean? Unfortunately, we do not really know how colour vision works, in what 'typical' ways spectral cues are employed. What we do know is that the colours we see depend upon the current ambient light plus the profile of wavelengths that specific materials are disposed to reflect. Further, we suspect that spectral signals are involved in just those visual tasks for which intensity cues prove inadequate. For example, it is often postulated that such cues are used to define equiluminescent borders, highlight the contrast between object and background, and to differentiate objects that are similar in all other respects (e.g. the ripe and unripe pear). (For a short explanation of colour pathways, see De Yoe and Van Essen, 1988; for a more thorough review of colour vision, see Gouras, 1984.) In other words, while we may think of the colour system as whatever neural machinery produces colour sensations, the colour system is more than that: it is that part(s) of the visual system that responds to, discriminates and utilizes spectral cues. It is this system, then, whatever it might be, that is activated when we see the film of the 'bat experience'.

Needless to say, a bat's colour sensations of acoustic stimuli would be quite another matter. Its sensations would not be tied to the ways in which external objects reflect ambient light nor would its sensations be a part of a system that uses the spectral composition of light for various information-processing tasks. The bat's colour sensations would be linked to properties of acoustic stimuli and to its auditory processes involved in spatial processing. As it turns out, although the bat film was presented as consisting of seemingly random coloured patches, I had in mind a specific process for the generation of those images. There was an informational relation between the properties of the visual image and those of the acoustic stimuli about which you, the 'viewer', were not told. That relation was as follows. First, the hue of the sensations (red, green, blue, etc.) encoded the frequency of the sound waves; second, the brightness of the colours gave the volume or intensity of the sound; and, third, the configuration of the patches showed, straightforwardly, the spatial properties of the sound waves. Finally, the film encoded the time delay of the echo or the bat's distance from surrounding objects. By making the coloured patches appear at different depths, spatial disparity mimicked a disparity in time – the amount of time it takes for the bat's outgoing cry to bounce off a distant object and return. The longer the delay between the cry and the echo, the further 'back' the coloured patches appeared in the 'visual' field. In this way, distance was represented by stereoscopic display.[2] Now, such an image of the sound field, in itself, would not buy the bat a sensory system for spatial perception. In order for the bat to perceive spatial relations in the world, something more would be needed: the visual images would have to be hooked up with various other neural processes 'further down the line' – with the bat's cortical pattern analysers that decode object shape, texture and identity, with the bat's vestibular and motor systems, and with, well, who knows what else? The fiction of the bat film, however, is that these colour sensations are what the bat experiences, *qualitatively* – a coloured image of the sound field, over time, as the bat pursues a mealworm.

One problem with the bat film now looks relatively clear: as a result of the differences between the human visual system and the bat auditory system, we cannot expect that by inducing colour sensations in ourselves we will understand the role that such sensations play in the bat's phenomenal world.[3] Because a 'sen-surround' film produces our visual experience through the usual means, we see the colours as we normally do, as the projection of moving coloured images upon a curved screen. Lacking the auditory/representational capacities of the bat, we do not experience the colours as does the bat, however that might be. All a film can show us are meaningless (albeit coloured!) visual events. Put another way, what the bat film seems to prove is that it is not for lack of the 'quality' of the bat's experience that his world eludes us. Even if, *ex hypothesi*, we were able to produce in ourselves the 'very feel' of the bat's experience, its 'qualitative' aspect, we would not understand the bat's point of view. Watching the swirl of colours, those sensations lack their proper representational content. We cannot expect to understand the bat's point of view, in other words, without access to both the representational and qualitative parts of its experience. And here we are given but one aspect, the phenomenological 'feel' of the bat's world.

Unfortunately, this way of putting things is not quite right, for it does not get to the root of the problem, does not fully explain why a film cannot give us the point of view of the bat. Let me try a different path. Both the description of the bat film as initially given and the conclusions drawn from it above presupposed that there could be a separation of the 'qualitative' and 'representational' aspects of phenomenal experience. 'What the bat hears is just like colour' the reader is told, 'except, of course, the colours mean something quite different. Imagine that!' This was how the thought experiment got off the ground. Yet sensible as that request might have seemed, we have no idea how to comply with it, what such a separation could be. As Daniel Dennett has often pointed out (see, for example, Dennett, 1988), what one is asked to imagine, what one can imagine and what one actually imagines are three distinct things. It is not clear that we do know how to separate our conscious experiences into two parts, the representational and qualitative aspects, or whether, indeed, this notion even makes sense. To illustrate this point, suppose that, instead of referring to the bat film, I had requested that you do the following:

> Open your eyes and look around your office (it's the end of term) – at the stacks of books and papers, at the piles of articles, unopened mail and ungraded papers. Note the way the scene looks to you, the inner phenomenology of the event. Now, a bat's consciousness is just like that – the feel of the scene is exactly the same – except, of course, all those visual sensations *mean* something very different to the bat. They represent quite different properties. Imagine that!

The problem is that you cannot imagine that, no matter how sincerely or hard you try. First, it would require that you 'strip away' the representational content of the entire office scene (say, by erasing the 'black lines' of the image, leaving only the 'crayoned' parts?). Then, by some other process, the intentional content of the bat's representations must be 'overlaid' upon the remaining bare sensory qualities (by a process akin to drawing in new lines or attaching new labels?). This, I contend, is not something we have any idea how to do: we do not know what the two 'parts' would be like, of and by themselves, so we have no inkling how to pull them apart or put them together. Our intuitions do not provide a concrete distinction between the qualitative and representational aspects of perceptions.

Still, you might well ask, why then, if there is no such distinction, did the bat example work at all? That is, in the bat film, we were asked to imagine meaningless coloured patches swirling across the screen – and we did. It also seemed perfectly reasonable to imagine that those colours played a representational role in the bat's experience, one that was different from the role they play in our conceptual scheme. But if there is no distinction between the

qualitative and representational parts of experience, how could this be so? Certainly it seemed to us that we could imagine such a distinction.

The answer here is that the description of the film was intentionally misleading: it was designed to play upon a common experience, that of seeing images or pictures we can not identify. Staring at an abstract painting perplexedly, we scan the blobs of colour for form – what could that possibly be a picture of? – when, suddenly, the figure of a man emerges. The apparently meaningless blobs of paint are transformed into a comprehensible image. These are the cases in which we legitimately regard content and 'mere colour' as distinct: at first the canvas contains only formless coloured blobs; after the 'aha!' experience, the painting has meaning – and this despite the fact that the canvas remains physically unchanged. It was this kind of event that set the stage for the original bat film. Given our familiarity with pictures and drawings, we tried to imagine a similar kind of thing – a film of 'meaningless' coloured shapes, non–intentional and non-representational sensory qualities, such that, if only we knew the proper 'squint' of the bat, those images would have content for us as well. We imagined, or at least we thought we could imagine, an unchanging substrate of pure sensation – a substrate analogous to the physical paint upon the canvas – onto which the bat's meaning could be affixed. The problem, however, is that our experience of abstract art does not provide a genuine example of what we need, the separation of content from 'mere colour'. Viewing an abstract painting does not involve an experience of a 'meaningless' image in the proper sense, that is, because the sudden emergence of a form in an abstract artwork is not the experience of having sensory stimuli, devoid of content, instantaneously gain representational properties. Even if we do not initially see the coloured shapes as the ghostly portrait of a man, we do see the colours as something – as coloured shapes upon a canvas, external to us, 3 ft dead ahead. The same is true for the patches of colour in the bat film. Perceiving (or imagining) moving coloured patches on a screen is an intentional – or at least, quasi-intentional – event, an experience of coloured patches as coloured patches. So when we imagined the bat film, we did not thereby imagine pure sensory qualities, colour qualia devoid of content. Our understanding of abstract art forms was misleading because it fostered the illusion that we could imagine exactly that.

Where does this leave us with respect to Nagel's original question and its intuitive basis? In questioning whether we could ever understand an alien organism's point of view, we intuitively construe this problem as analogous to the everyday task of understanding the phenomenal experiences of each other. Here, because our own difficulties turn around individual sensations, around the 'feel' of sensory events – the pain of a migraine headache, the azure blue of the Mediterranean, the 'essence' of flamingo pink – we infer that the main stumbling block to understanding an alien creature must be the inaccessibility of those qualia. We treat a conscious experience, in other words, as a mere collection of qualia, as a bunch of individual sense data that have somehow come together to form a phenomenological whole. (Certainly, this is the route that most analytic philosophical debates have also taken. In the 'inverted spectrum' problem, for example, the question is asked whether it would be possible for two people to have exactly the same neural structures and functions and yet have their colour experiences be 'spectral inversions', one of the other. Could you, my neurological equivalent, see the sky as red even though I see it as blue? In the 'absent qualia' problem (Block, 1978), the question is whether an artificial system functionally identical to one's own brain could be entirely devoid of qualitative experience. If, given a Turing-machine table that described the functional states of my brain, the entire population of China could be talked into instantiating, for one hour, the state types specified by that table, would my aches, tickles and pains be somehow 'experienced' (collectively?) by all the citizens of China? These are the kinds of questions – questions phrased in terms of individual sensations – that are currently asked.)

What is overlooked by the intuitive construal of the problem are the following two points. First, because we are able to individuate, identify and catalogue some of our phenomenological experiences and to converse with other people about such perceptual experiences as 'that very colour' (referring, say, to the intense blue-green of the Mediterranean), it does not follow that these sensations come to exist *in vacuo*. This 'isolation' of those sensations (whether as a result of some internal process of individuation or merely in virtue of linguistic convention) does not thereby produce sensations that stand apart from our representational/conceptual schemes. What the intuitive view conflates, in other words, is an ability to refer to certain parts of conscious intentional experience with an ability to pick out its purely qualitative aspects. Isolation does not distil qualia from content. So, whatever the root of our everyday problems in communication, it is not the intrinsic nature of sensations *per se* that makes for trouble – or, rather, there is no reason to think that this is the case given our communicative problems. If our utterances do not refer to pure sensation, one sees that the problems of communicating our phenomenological experience are equally a problem about representational states.

Second, a point of view, as we know from our own – paradigmatic – case, is not a jumble of qualia. In the normal non-pathological subject, consciousness is systematic, representational and intentional (e.g. we represent objects as being a certain way or of a certain type). Moreover, such properties are not 'optional' parts of our conscious experience, merely accidental or inconsequential aspects, if they can be considered 'parts' at all. Rather, these properties are constitutive of a point of view. That we experience the world in any way at all – that it is like anything to be me – is made possible by exactly these properties. So, given that our own phenomenal experience is the starting point for an explanation of the very notion of a point of view, and that our own experience is not a mere collection of qualia, we must assume that the same holds for the bat. If there is anything it is like to be bat, we have no reason to think – indeed, there is no sense to the suggestion – that that bat's experience is but a collection of pure qualia.

The mistake of the intuitive view, then, was first to think that our problem of communication was one about pure qualitative states, and then, second, to import this interpretation of the problem into the task of understanding an alien point of view. If we construe our communicative failures to hinge upon pure qualitative states of which the speakers do not have a common experience, then what we face in understanding a foreign phenomenology is simply 'much more of the same' – for the bat will have more and more purely qualitative states of which we ourselves have had no experience. By misconstruing the nature of an interpersonal problem, the puzzle about another creature's point of view becomes a problem about pure qualia.

The upshot of the bat film, then, is this. Nagel has claimed that we will never understand the point of view of an alien creature. This is a claim that our intuitions support with a nod towards 'that something', pure phenomenal experience, which cannot be known merely by description, without personal experience. But if introspection does not yield any clear distinction between the representational and qualitative properties of experience, then we do not know, *a priori*, what insights or even what kinds of insights will result from empirical investigation. Certainly we cannot confidently declare that science must fail to unearth 'that something', for we have no clear idea to what this amounts; nor can one say what the scientific approach will necessarily leave out, if it must leave out anything at all. This gives us, I think, good reason to continue on with our empirical investigations of mental representation – to look towards the disciplines of neurophysiology, psychology and artificial intelligence – without undue pessimism about the relevance of their experimental results.

Ourselves as Subject

One consequence of tying together sensation and representational experience is that the nature of our own subjective experience is opened to investigation (Sellars, 1963; Dennett, 1978a; Churchland, 1983). It is as legitimate a subject of inquiry as the experience of other creatures. Because the questions about phenomenology are no longer focused on the intrinsic quality of particular sensations but on a phenomenology as a whole – complete with its representational/intentional nature – our ignorance extends to ourselves as well. We, as the 'owners' of our point of view, do not thereby understand its representational character. Hence, our study of representational systems is also an investigation into our own point of view.

This consequence is, I suspect, somewhat counter-intuitive. If anyone knows about my subjective experience, it is certainly me, or at least that is what we have always thought about the matter. By way of lending some small amount of plausibility to this result, then, I want to end this chapter by going back to the example at the beginning, that of the eagle. What did learning a simple anatomical fact about the eagle, about the foveation of the eye, tell us about that creature's experience? More importantly, how would a fact about an eagle nudge our sense of self, reflect upon the human experience?

In learning that the eye of the eagle has two separate foveal regions, it suddenly seemed clear that the experience of the eagle must be different from our own. On the other hand, when I tried to imagine *how* the experience of an eagle would differ from my own, I immediately adopted a hypothesis that incorporated my own visual system into the experience. I wondered, that is, whether being an eagle might not be akin to the experience I would have while wearing strange quadra-focals – whether it wouldn't be like shifting my own gaze from lens to lens sequentially. In essence, I incorporated my own foveal field into the experience of being an eagle. (This would give me, in effect, eight different levels of visual acuity: four lenses imposed upon my foveal and non-foveal regions.) Of course, nothing we know about the visual system of the bird of prey constrains its visual 'attention' in a similar way. Although my foveae must move from lens to lens sequentially, the bird need not have any analogous 'inner' eye that receives, serially, the information from the two foveae and the horizontal band. Because there are parallel lines from all regions of the retina, there is no reason why the brain must process the information sequentially – no reason why, say, the eagle must first attend to the left, then forward, then to the horizon just as I would. The eagle might 'attend' simultaneously to all this information at once, no matter how this might conflict with our intuitive notion of visual attention. This is a possibility that the anatomical data reveals.

Note that once we see how a notion of 'foveal' processing has been misapplied to the eagle's point of view, it is an interesting question whether or not we have also 'moved the eye inward' not merely in thinking about the eagle, but alas in thinking about ourselves. Here, I am referring to the many models of conscious attention that utilize, in one form or another, the 'spotlight' metaphor: the 'inner eye' of consciousness shifts like a searchlight from one neural event to another, successively attending to different mental events. This, too, is a 'foveal' theory of attention, not of another organism's consciousness but of our own. We apply the foveal metaphor to our conscious experience as a whole. Certainly, this is a model with intuitive plausibility. Something about it seems just right. The question that the eagle's eye raises, however, is about the basis of this appeal. Is it appealing because this is, in fact, how our inner experience is, or does it seem right just because the foveated nature of our visual experience colours our understanding of conscious attentive processes as a whole?

First, the former alternative could be true. The spotlight theory might seem plausible because, on looking inwardly at ourselves, we can see by introspection that our consciousness is sequentially focused on single events. That is, the introspective evidence coheres with the metaphor. But is this really so? Recall what it is like to struggle through a recalcitrant screen door weighed down by several bags of groceries. First, you juggle the groceries and grasp the door handle; then you feel a mosquito land on your ankle; then you hear the creaking door hinge and the rip of a paper bag; then the

mosquito makes a stab with his proboscis; then you loose your grip on the handle; then the screen slams shut on your shin; then a tin can bounces off your thigh . . . Somehow, this strictly sequential narrative does not quite capture the experience, even if it does record the objective order of the external events. The very problem with such experiences is that 'everything happens at once'. In the midst of the calamity, what happens first – the bag ripping or the mosquito biting or the screen door slamming – is not always clear. On the basis of experience alone, there is no distinct ordering of all of the events, no clear sequence of this event, then this one, then this and finally that.

Perhaps, then, the explanation goes the other way about: perhaps the searchlight metaphor, combined with our story-telling practices and our understanding of the relevant causal chain of events, confer order upon the conscious events only in retrospect. What I am suggesting is that the spotlight metaphor may be adopted just because (a) we are foveated animals and (b) we do not actually perceive any firm order in the events (i.e. such events are not 'tagged' for time). Because we are such strongly visual organisms and because eye movements are required for our perception of the world, the metaphor seems plausible. Needing an explanation, we mistake our intuitive grasp of the visual perception of external events for an accurate description of internal attentional processes. We co-opt the visual notions of 'searching', 'focusing' and 'watching' and apply them to all of conscious experience. This, I think, is possible. What the eye of the eagle should make us wonder is whether our conception of ourselves might not be 'tainted' with the same foveal metaphors we naturally apply to other creatures.

The above example is not meant as a serious criticism of spotlight theories of conscious attention. Rather, it is given as a suggestive example of how it could come about that we are mistaken about our own inner events – how the way our own attentional mechanisms seem to us could diverge from how in fact they are. It offers a small glimpse of the ways a possible reconception of ourselves, and our point of view, could come about in the light of physiological/computational discoveries.

Still, the central idea of this chapter has been that we do not know what science will explain, just because we lack a firm grasp on the subject matter: the nature of conscious events. If so, we are in a funny position. We will know what science can tell us only after it has done so. Hence, only suggestive examples are now possible. What we can provide, however, are good reasons to wait – to see what science will do. In effect, this is what I have been attempting to show in this chapter.

Acknowledgements

This chapter began as an introduction to paper, 'What is it Like to be Boring and Myopic?' (Akins, 1993) where it served, in a much abbreviated form, to motivate the neuroscientific approach to the problem of consciousness used there. An earlier version appeared under the title 'Science and Our Inner Lives: Birds of Prey, Bats and the Common (Featherless) Bi-ped' in a collection edited by Marc Beckoff and Dale Jamieson (1990). For their generous comments on and discussion of the manuscript, I would like to thank Marc Beckoff, Daniel C. Dennett, Dale Jamieson, Joseph Malpeli, Wright Neely, Brian C. Smith, Tony Stone, Tom Stoneham and Mary Windham. I would also like to thank Martin Davies for his extensive comments on the final draft.

Notes

1 That someone being Jeremy Butterfield.
2 This way of generating the film was given only for the sake of example, not because I think that this is what a bat's experience is really like. That is, assuming that a bat does have a point of view (and I doubt that it has), the film represents the properties of the sound field before the sound waves are transduced, processed and filtered by the basilar membrane, midbrain and auditory cortex of the bat. At the level of the auditory cortex (surely the first neural level at which conscious experience would be possible), the informational characteristics of the signal have been significantly changed.
3 It is an interesting question, however, whether, given the addition of dopplershift or velocity information to the visual display, our own visual systems could act as a spatial pattern analyser of some sort – that is, whether if we, given the intellectual knowledge of how the image is produced, were to look at the screen we could learn to use that information to guide our actions, say to walk around a room filled with objects.