

VOYAGES OF DISCOVERY

Neurophysiologists, most of them, are still under the influence of dualism, however much they deny philosophizing. They still assume that the brain is the seat of the mind. To say, in modern parlance, that it is a computer with a program, either inherited or acquired, that plans a voluntary action and then commands the muscles to move is only a little better than Descartes's theory, for to say this is still to remain confined within the doctrine of responses.

—J. J. Gibson

In this chapter I tell the story of Hubel and Wiesel's Nobel Prize-winning research into vision in mammals. The work rests, I show, on an untenable conception of vision and other mental powers as computational processes taking place in the brain. The main problem with the computational theory of mind is that it supposes, mistakenly, that mind arises out of events in the head. The legacy of Hubel and Wiesel's research must be called into question.

The Visual Brain in Action

In 1981, David Hubel and Torsten Wiesel were awarded the Nobel Prize for research on the neurophysiology of vision, research

they had conducted at Johns Hopkins, and then at Harvard, from the late 1950s until about 1980. Hubel and Wiesel's research, and the fact that it earned them the highest acclaim from the scientific establishment, is an important landmark in the science of consciousness. Seeing, after all, is in the first instance a mode of animal consciousness. In fact, for humans at least, seeing plays a colossal role in our conscious lives. The world is open to our visual inspection and we rely on seeing to get what we want and find our way around. But more than that, for us human animals our world is a visual world. It is a world full of profiles and colors and vistas. The visual character of objects shapes how we conceive of them: we think of them as having, for example, fronts and backs and hidden aspects. Think how difficult it would be to frame an understanding of what is going on around us—or at least to frame an understanding like that which we enjoy—if we couldn't see.

It is sometimes said that we know more about how the brain enables us to see than we do about any other mental function of the brain. When people say this they usually have the work of Hubel and Wiesel in mind. For most of history it has been impossible to study the workings of the brain of a living human being or animal. How would you do it? The brain is not open to view; it is hidden in a case of bone. And even if it were not—even if the skull were transparent—the brain's functional organization is obscure and its complexity makes it explanatorily opaque.

Hubel and Wiesel's importance is that they seemed to find a way to exhibit the brain's workings in a way that made what it was doing intelligible: how the brain might be achieving the function of making us visually conscious. To an astonishing degree their work was and remains the standard by which the field measures itself. As I will now explain, the shortcomings of Hubel and Wiesel's approach remain, even now, the shortcomings of the field of research into the neural basis of consciousness.

The Basic Project

Hubel and Wiesel's story, the journey they took, is fascinating and instructive. Let's begin at the beginning. First, Hubel and Wiesel poked fine microelectrodes into the visual cortices of cats and monkeys in order to record the electrical behavior of single cells. Doing so inflicted damage on the animals—after all, the electrode had to cut a path through tissue. But damage was isolated and this procedure seemed to allow, at least for a time, the investigation of the more or less normal behavior of individual cells.

Hubel and Wiesel were not the first to record in this way from the cortex: Vernon Mountcastle, at Johns Hopkins University, had recorded from the somatosensory cortex. Others, notably Sir John Eccles, the great Australian physiologist, had pioneered the technique of single-unit (cell) recordings in the spinal cord. One contemporary of Hubel and Wiesel's, Jerome Lettvin, at MIT, remarked that he rejoiced that Eccles had rescued neurophysiology from "Sherringtonian ooze." What Lettvin seems to have meant was that it was Eccles who had first found a way to translate the macroscopic neurophysiology of Sir John Scott Sherrington to the microscopic level.

Stephen Kuffler, who was Hubel and Wiesel's mentor at Johns Hopkins, and also Horace Barlow, at Cambridge University, who was their rough contemporary, had by the mid-1950s made important discoveries about the behavior of retinal cells. The receptive field for a visual cell is the area of the retina whose stimulation causes the cell to alter its firing rate. (Alternatively, you can think of the cell's receptive field as the region in space in front of the animal to which a cell is responsive.) Kuffler found that retinal ganglion cells had receptive fields consisting, in effect, of concentric circles. For such "on-center" cells, a spot falling in the central region of the receptive field would activate the cell, whereas a ring-shaped annulus falling outside the cen-

tral region would inhibit the cell's firing. A diffuse light falling equally over the whole receptive field would produce a weaker reaction than a spot falling only on the center. Off-center cells had it the other way around.

Hubel and Wiesel were impressed with Kuffler's discovery, and from the start it was clear what they would try to do: "The strategy . . . seemed obvious," Hubel wrote. "Torsten and I would simply extend Stephen Kuffler's work to the brain. We would record from geniculate cells and cortical cells, map receptive fields with small spots, and look for any further processing of visual information." Others had tried to do this, but without any significant success. The problem, it turned out, was figuring out what sorts of stimuli would produce activation in cortical cells; in other words, the problem was getting the cortical cells to respond at all. "The cells simply would not respond to our spots and annuli," Hubel lamented. They did crack the problem eventually; Hubel and Wiesel were the first scientists who figured out how to make cells in the visual cortex talk, as it is sometimes put. Their first discovery came about by accident. They were trying to find ways to stimulate a cortical cell by using slides to project spots onto a screen in front of the animal. No matter where they projected the spots, they couldn't get any response in the cell they were measuring from. Hubel explained:

Then gradually we began to elicit some vague and inconsistent responses by stimulating somewhere in the mid-periphery of the retina. We were inserting the glass slide with its black spot into the slot of the ophthalmoscope when suddenly over the audiometer the cell went off like a machine gun. After some fussing and fiddling we found out what was happening. The response had nothing to do with the black dot. As the glass slide was inserted its edge was casting a faint but sharp shadow, a straight dark line on a light background. That was what the cell wanted,

and it wanted it, moreover, in just one narrow range of orientations.

What Hubel and Wiesel had discovered was a cell whose response was triggered by lines at a certain orientation. After this initial discovery, progress was steady, if arduous. Hubel and Wiesel found classes of cells in the cat visual cortex whose receptive fields departed strikingly from those found in the retina or in the geniculate (a thalamic way station between the retina and the cortex). For example, they found cells that had receptive fields with a similar sort of antagonistic organization to that found in Kuffler's center/surround ganglion cells but without the circular symmetry characteristic of cells in the retina. The optimal stimuli for these cells were stationary lines and slits at particular positions and with fairly precise orientations. They called these cells "simple." They also found cells that were like simple cells insofar as they responded best to lines or edges of given orientations but, unlike simple cells, were insensitive to the position of the line within the receptive field. Hubel wrote that the behavior of these cells "can most easily be explained by supposing that the complex cells receive inputs from many simple cells, all of whose receptive fields have the same orientation but differ slightly in position." What Hubel and Wiesel came to believe is that the network of cells as a whole is organized hierarchically; that is to say, complex cells are driven by networks of simple cells.

This was just the beginning. Among the highlights of Hubel and Wiesel's nearly twenty-five-year research collaboration can be counted the finding that complex cells give an especially strong response when a line is swept across the field, with some cells firing better to one direction of movement than another, and also the discovery of even more specialized "hypercomplex" cells. These specialized cells were frequently said to be selective for orientations and directions of movement.

Hubel and Wiesel made significant strides in characterizing

tral region would inhibit the cell's firing. A diffuse light falling equally over the whole receptive field would produce a weaker reaction than a spot falling only on the center. Off-center cells had it the other way around.

Hubel and Wiesel were impressed with Kuffler's discovery, and from the start it was clear what they would try to do: "The strategy . . . seemed obvious," Hubel wrote. "Torsten and I would simply extend Stephen Kuffler's work to the brain. We would record from geniculate cells and cortical cells, map receptive fields with small spots, and look for any further processing of visual information." Others had tried to do this, but without any significant success. The problem, it turned out, was figuring out what sorts of stimuli would produce activation in cortical cells; in other words, the problem was getting the cortical cells to respond at all. "The cells simply would not respond to our spots and annuli," Hubel lamented. They did crack the problem eventually; Hubel and Wiesel were the first scientists who figured out how to make cells in the visual cortex talk, as it is sometimes put

Their first discovery came about by accident. They were trying to find ways to stimulate a cortical cell by using slides to project spots onto a screen in front of the animal. No matter where they projected the spots, they couldn't get any response in the cell they were measuring from. Hubel explained:

Then gradually we began to elicit some vague and inconsistent responses by stimulating somewhere in the mid-periphery of the retina. We were inserting the glass slide with its black spot into the slot of the ophthalmoscope when suddenly over the audiometer the cell went off like a machine gun. After some fussing and fiddling we found out what was happening. The response had nothing to do with the black dot. As the glass slide was inserted its edge was casting a faint but sharp shadow, a straight dark line on a light background. That was what the cell wanted,

and it wanted it, moreover, in just one narrow range of orientations.

What Hubel and Wiesel had discovered was a cell whose response was triggered by lines at a certain orientation. After this initial discovery, progress was steady, if arduous. Hubel and Wiesel found classes of cells in the cat visual cortex whose receptive fields departed strikingly from those found in the retina or in the geniculate (a thalamic way station between the retina and the cortex). For example, they found cells that had receptive fields with a similar sort of antagonistic organization to that found in Kuffler's center/surround ganglion cells but without the circular symmetry characteristic of cells in the retina. The optimal stimuli for these cells were stationary lines and slits at particular positions and with fairly precise orientations. They called these cells "simple." They also found cells that were like simple cells insofar as they responded best to lines or edges of given orientations but, unlike simple cells, were insensitive to the position of the line within the receptive field. Hubel wrote that the behavior of these cells "can most easily be explained by supposing that the complex cells receive inputs from many simple cells, all of whose receptive fields have the same orientation but differ slightly in position." What Hubel and Wiesel came to believe is that the network of cells as a whole is organized hierarchically; that is to say, complex cells are driven by networks of simple cells.

This was just the beginning. Among the highlights of Hubel and Wiesel's nearly twenty-five-year research collaboration can be counted the finding that complex cells give an especially strong response when a line is swept across the field, with some cells firing better to one direction of movement than another, and also the discovery of even more specialized "hypercomplex" cells. These specialized cells were frequently said to be selective for orientations and directions of movement.

Hubel and Wiesel made significant strides in characterizing

the "functional architecture" of the visual cortex, finding, for example, that columns of cells with related receptive field properties formed functional units. Moreover, orientation columns were found to exist that are in effect, as Hubel put it, like "a little machine that takes care of contours in a certain orientation in a certain part of the visual field."

And then there are experiments on the consequences, for cortical development, of depriving newborn cats and monkeys of the use of their eyes by sewing their lids shut. Hubel and Wiesel demonstrated that deprivation during this critical period caused an irreversible lack of connectivity in the cortex resulting in permanent blindness. They showed that the ability to see requires experience. If prevented from seeing during a critical period, animals will never see.

Christopher Columbus and the Brain

Hubel and Wiesel's findings are impressive. These are hard facts, and they speak for themselves. That seems to have been Hubel and Wiesel's view. Hubel wrote: "Almost absent from our way of working and thinking were hypotheses, at least explicit ones. We regarded our work as mainly exploratory, and although some experiments were done to answer specific questions, most were done in the spirit of Columbus crossing the Atlantic to see what he would find." And he observed: "It is hard, now, to think back and realize just how free we were from any idea of what cortical cells might be doing in an animal's daily life."

What a strange and remarkable thing for Hubel to say! Christopher Columbus didn't just set off across the Atlantic to see what he would find. He had very precise and, as we now know, erroneous ideas about what he would find. But the storied explorer aside, it is impossible to take seriously the claim that

Hubel and Wiesel were not guided by theory and responsive to its demands. Indeed, how could they *not* have been? After all, there are billions of cells in the brain, and they are massively interconnected. To form any conception whatsoever of what individual cells are contributing to the brain's functioning, you need to have a tolerably clear idea in advance of what the brain is doing. And indeed, Hubel and Wiesel did have such a guiding conception.

In *Eye, Brain, and Vision*, published in 1995, Hubel wrote: "We know reasonably well what it [the visual cortex] is 'for,' which is to say that we know what its nerve cells are doing most of the time in a person's everyday life and roughly what it contributes to the analysis of visual information." And he went on to add: "This state of knowledge is quite recent, and I can well remember, in the 1950s, looking at a microscopic slide of visual cortex, showing the millions of cells packed like eggs in a crate, and wondering what they all could conceivably be doing, and whether one would ever be able to find out."

Back in 1958, when Hubel and Wiesel set off on their journey of discovery, no one knew how neurons functioned so as to contribute to the analysis of visual information. In that sense it is true that it was not known back then what the visual cortex is "for." But that the visual cortex was in the business of analyzing visual information (as Hubel put it) and that, therefore, individual neurons must somehow, some way, be making a contribution to that was something that Hubel and Wiesel knew from the moment they set sail. Or rather, it was something they took for granted and assumed. Consider again Hubel's remark, first cited in the last section, now with italics added: "The strategy... seemed obvious. Torsten and I would simply extend Stephen Kuffler's work to the brain. We would record from geniculate cells and cortical cells, map receptive fields with small spots, and *look for any further processing of visual information.*"

And they were not alone in taking an information-processing conception of the brain for granted. By the late 1950s it was common belief among neuroscientists that vision presented the brain with a problem in information processing, and that the parts of the brain dedicated to vision could be thought of as systems of networks or circuits or, as Hubel and Wiesel sometimes put it, machines for "transforming information" represented in one system of neurons into progressively more refined and complex representations of what is seen. For Hubel and Wiesel, the visual system consisted of cells whose receptive field properties made them, in effect, symbols for features such as edges, orientations, directions of movement, and color. Cells were understood to be specialized in order to be able to "stand for" and thus represent features. This application of information theory to the brain was not new when Hubel and Wiesel set to work. Santiago Ramón y Cajal's student Rafael Lorente de N6 had represented neural relations as networks already in the thirties, and his treatment had a direct influence on the work of Warren McCulloch, Walter Pitts, and, through them, John von Neumann. (Walter Freeman, the neuroscientist, likes to say that Lorente de N6 is in a way the godfather of the digital computer.) Interestingly, Claude Shannon, one of the developers of the mathematical theory of information, was skeptical that the brain was an information processor; he believed that processing information required a transmitter, a receiver, and an agreed-upon code, none of which is found in the brain. But Shannon's skepticism didn't do much to dampen the general enthusiasm for the new approach. And so, twenty-five years after the beginning of their joint collaboration, Hubel and Wiesel were awarded the Nobel Prize "for their discoveries concerning information processing in the visual system."

Hubel and Wiesel set out from the start to understand how the behavior of individual cells, and their organization into

larger assemblies, could accomplish the information-analysis task that is vision. They took for granted that vision was a process of analysis of information. It is remarkable that their landmark investigations into the biology of vision take as their starting point a startlingly nonbiological engineering conception of what seeing is.

The Computer Model of the Mind

David Marr, whose landmark book *Vision* was published in 1982, the year after Hubel and Wiesel won the Nobel Prize, made explicit the conception of vision on which Hubel and Wiesel implicitly relied. Marr pronounced that vision is an information-analysis process carried out in the brain. This, of course, was how Hubel and Wiesel had understood vision all along. Vision is the process of discovering how things are in the scene from images in the eyes. That is, it is a process of extracting a representation of what is where in the scene from information about the character of light arrayed across the skin of receptors in the eyes.

As we have noted, the use of information-theoretic ideas to understand what the brain is doing was already well entrenched by the late fifties. As far back as the nineteenth century, Helmholtz had proposed that perception is inferential: the brain constructs and tests hypotheses about what sort of events in the world are producing these impressions. And indeed, as we saw in the last chapter, this has been a guiding idea in the study of vision over the course of the last century.

What was new in Marr's work—here he made an advance over Hubel and Wiesel—was the attention to and conceptual clarity regarding theory that he brought to the table. Marr wrote that "trying to understand vision by studying only neurons is like trying to understand bird flight by studying only feathers. It just

cannot be done." You need a theoretical conception of what neurons (or feathers) are doing just in order to decide which facts are even relevant. That is, you need to characterize what the system is doing in a way that is more abstract. And this is not because of anything peculiar to vision but because of the particular explanatory challenges we face when we want to understand an information-processing mechanism.

Take a simple case. You can't understand how a particular cash register works if you don't understand what it is for, namely, adding up numbers to keep track of balances due. When it is clear that that's what the machine is doing, then you can reasonably ask: How does it manage to do this? And so you will have to investigate the different ways this machine, or any machine, might be organized so as to add up figures.

There are lots of different procedures or recipes for doing this—what mathematicians call algorithms. When you select an algorithm, you select, in effect, a way of representing a problem and a way of solving it. For example, we are all familiar with algorithms for doing addition that involve the use of a pencil, paper, and Arabic notation. You don't need to understand addition in any deep sense in order to add up numbers. Now, there are different algorithms for adding: you'd follow different procedures if you were representing numbers in Roman notation or in binary notation. Likewise, there are lots of different kinds of physical mechanisms that can be used to carry out the function of addition. You might use pencil and paper, or an abacus, or a mechanical cash register, or a digital computer. In your effort to understand how a given machine works, you need to grasp and answer three questions: First, what function is it computing? Second, what algorithms or rules is it using to carry out that function? Third, how are those algorithms implemented physically in the mechanism?

The beauty of this approach is that it enables the study of an

information-processing mechanism to move forward even though the physics or electronics or physiology of the mechanism may be unknown. If vision is the process of producing a representation of the scene from information about the wavelength or intensity of points of light striking the eyes, then we can begin to investigate what kinds of rules would enable that kind of analysis of visual information even before we understand much about the behavior of cells in the eyes and the way they are linked together in networks. The information-processing approach to the mind, and to vision in particular, enables one to appreciate both that the processes are carried out in a physical medium (the brain, a computer, whatever) and that the processes are not themselves intrinsically physical. They are information-theoretic, or computational.

Again, we come up against the irony that vision is made amenable to neurophysiological investigation only at the price of conceptualizing vision as, in and of itself, a nonbiological (that is to say, a computational) process that just happens, in humans, to be realized in the brain. The fact that we can see thanks only to the workings of our wet, sticky, meat-slab brains doesn't make seeing an intrinsically neuronal activity any more than chess is. To understand how brains play chess, you first need to understand chess and the distinct problems it presents. And, crucially, you don't need to understand how brains work or how computers are electrically engineered to understand that. Chess is only played by systems (people and machines) made out of atoms and electrons. But chess isn't a phenomenon that can be understood at that level. And the same is so for vision. To understand how the brain functions to enable us to see, according to the information-processing perspective, you must understand vision as the sort of process that might unfold just as easily in a computer.

Is the Brain an Information Processor, Really?

Marr's idea, and Hubel and Wiesel's, is that the visual system—the vision-dedicated parts of the brain—performs an information-processing task: it extracts information about the environment from the retinal image, thus constructing an internal representation of that environment. For example, the brain notices sharp discontinuities in the intensity of light at different points; in its internal representation of the scene, it labels these places "edges." That's just what seeing is: a process whereby the brain takes patterns of light on the retina and transforms them into a representation of what is where in the scene before the eyes.

The information-processing approach to the brain and vision has been entrenched in science for almost a century. You might turn on the radio any day of the week and hear an author blandly state, as a matter of established fact, that language is "processed" in the left hemisphere of the brain or that it is the neocortex that computes higher cognitive functions. And we are not in the least nonplussed to learn that Marr, Hubel and Wiesel, and others hold that seeing is a neural process in which information is extracted by the visual system from the retinal image.

But is the brain an information processor, really? There is an obvious reason to question this conclusion. Consider: we know what it means to say that a detective, for example, extracts information about an intruder from a footprint, or that an oceanographer gathers information about a prehistoric climate by studying fossils of unicellular organisms that she dredges up from today's ocean floor. These are nice examples of "extracting information" about one thing from another. The explanation of the fact that the footprint and the fossils contain information about the intruder and the climate, respectively, is the further fact that there is a definite causal relationship between the character of the intruder and the properties of the footprints, or between the cli-

mate millions of years ago and the fossil chemistry of foraminifera today. And what makes it the case that the detective and the oceanographer can extract this information is that they are each armed with knowledge of the way in which what they have access to now (the footprint, the fossils) was shaped by what they want to learn.

Things are different, though, when it comes to the brain and the retinal image. No doubt the retinal image is rich in information about the scene before the eyes; after all, there are reliable and well-understood mechanisms whereby the former is brought into being by the latter. Presumably, then, a suitably placed scientist would be able to extract that information. But the brain is no scientist or detective; it doesn't know anything and it has no eyes to examine the retinal image. It has no capacity to make inferences about anything, let alone inferences about the remote environmental causes of the observable state of the retina. How, then, are we to make sense of the idea that the brain is an information-processing device?

There is a danger of vacuity in this "computer model of the mind." Our goal is to understand the biological basis of mind. It is hard to see how we succeed in making a contribution to the attainment of this goal if we suppose that our own mental powers are to be explained with reference to the cognitive powers of the brain. We—adult humans and other animals—think; we see, we feel, we judge, we infer. It's working in a big, plain circle to say that what makes it possible for us to do all that—what explains these prodigious powers of mind—is the fact that our brains, like wily scientists, are able to figure out the distal causes of the retinal image. For that just takes for granted the nature of mental powers without explaining them. Is cognitive science guilty in this way of reasoning as if there were mind-possessing agencies (hominunculi) at work inside us?

The Computational Brain

You may think that the existence of the digital computer—I mean the ubiquitous, everyday consumer appliance—provides proof positive that a mere mechanism such as the brain can process information. Computers, after all, perform calculations; they render three-dimensional models from line drawings. Computers correct spelling and play chess, and they do so, we know, without magic or the use of little guys inside them. What better grounds could there be for thinking we should take seriously the thought that brains are, in effect, organic computers? Whatever mystery might seem to attach to thinking of the brain as an inferring, reasoning extractor of information shivers away when we consider that even much simpler human-made artifacts like computers are capable in this way of thinking.

Some problems admit of mechanical solutions. If you want to know how many people are in the room, you can count them. Coming up with the solution requires no more than an ability to add one again and again. Likewise, you don't need to understand long division in order to find the answers to long-division problems. You just need to be careful. You were taught a decision procedure in school, one that makes use of writing, the Arabic system of notation, and the fact that you know how to divide, add, subtract, and multiply very small numbers. Any sufficiently careful idiot can do it. A machine can do it. In this same way, you don't need to be a great chef to follow a great chef's recipe, and you don't need to be able to grasp the vast number of combinations and permutations possible on a Rubik's Cube to learn the tricks that allow you to "solve" it in seconds.

An algorithm is a recipe or procedure for solving a problem. It is, as it were, a program that enables one (a child, an idiot, a machine) to reach a desired conclusion in a finite number of steps. Some problems can be solved by algorithms; some cannot.

There is no general procedure for deciding, for any given puzzle or problem, whether it is or is not "decidable" (as mathematicians say) by purely "mechanical," formal methods. It has been shown, however, that any problem that is mechanically (or "effectively") decidable can be computed by any of a class of formal systems. A digital computer, as we know it today, is one such physically realized formal system.

But it would be a mistake to think that these findings in the mathematics of computability—or that achievements in the domain of computer engineering—prove that our brains are, in effect, computers. For this claim is founded on a mistake. No computer actually performs a calculation, not even a simple one. Granted, following a recipe blindly and without comprehension is one way to find the answer to a puzzle. But, crucially, understanding a problem or a computation does not consist in merely following a rule blindly. As a little reflection on one's school days will make evident, there's all the difference in the world between understanding the solution to a problem and getting a good score on the test because you have memorized a recipe for doing so. Computers may generate an answer, but insofar as they do so by following rules blindly, they do so with no understanding.

But more important, computers don't even follow rules blindly. They don't follow recipes. Just as a wristwatch doesn't know what time it is even though we use it to keep track of the time, so the computer doesn't understand the operations that we perform with it. We think with computers, but computers don't think: they are tools. If computers are information processors, then they are information processors the way watches are. And that fact does not help us understand the powers of human cognition.

The Mind Is Not in the Head

Now, the fact that computers don't think may be a good reason to hold that brains don't think *because* they are computers. The philosopher John Searle, my colleague at UC Berkeley, has made this point persuasively. This leads Searle to go on to claim that consciousness and cognition arise from the intrinsic nature of neural activity itself. They are "caused by and realized in the human brain." Computers solve problems and represent the world only derivatively, thanks to the fact that we treat them as if they do. But the brain's powers are not derivative; they are original. The brain thinks and represents.

But this is exactly the wrong conclusion to draw from the fact that brains don't think by computing: they don't, but not because they think some other way. Brains don't think. The idea that a brain could represent the world on its own doesn't make any more sense than the idea that mere marks on paper could signify all on their own (that is, independently of the larger social practice of reading and writing). The world shows up for us thanks to our interaction with it. It is not made in the brain or by the brain. It is there for us and we have access to it. What makes it the case that my thoughts are directed to this task (playing chess, say) or to this object (a glass of water, for example) is not the intrinsic nature of my internal computational states. I agree with Searle on this score. But that's because what gives my thoughts their content is my involvement with the world. On no account does my interior makeup suffice all on its own to give meaning and reference to my mental states. Meaning is not intrinsic, as the philosopher Daniel Dennett has rightly argued; it is not internal. Meaning is relational. And the relation itself thanks to which our thoughts and ideas and images are directed to events, people, and problems in the world is the fact of our being embedded in and our dynamic interaction with the things around us. The world is our ground; the world provides meaning.

The limitations of the computer model of the mind are the limitations of any approach, to mind that restricts itself to the internal states of individuals. Cognitive science sought to reveal how the brain can be a subject of thought by supposing that it is, in effect, a kind of digital computer. But what now emerges is that computers cannot think (or see, or play chess), and for the very same reasons that brains can't.

The central claim of this book is that the brain is not, on its own, a source of experience or cognition. Experience and cognition are not bodily by-products. What gives the living animal's states their significance is the animal's dynamic engagement with the world around it.

The Mind-Body Problem for Robots

The movie *Blade Runner*, discussed briefly in Chapter 2, makes vivid what seems to be incontrovertible: there is no principled basis for denying the rebellious worker slaves the respect and consideration we believe to be required of us in relation to human beings. Certainly, no facts about what is going on inside replicants justify such a denial. Replicants are made, yes. But then, in a sense, so are we. And yes, they lack a certain autonomy; but then we, too, lack autonomy. Far into our adulthood, we depend on parents, family, friends, society; *just to survive*. Granted, replicants lack biological innards; they are not composed of the same stuff as we are. But that's just the point: there is no necessary connection between what we are and what we are made out of. It would be nothing but prejudice to insist that there is such a connection.

Of course there are all sorts of practical reasons to think that we need brains like ours to have minds like ours. The technology that would support artificial minds is a long way off. The point is that it is not as though we understand so well how we work that

we can rule out, in advance of case-by-case consideration, whether we will one day discover, or even learn how to manufacture, different kinds of minds.

A curious upshot of this line of argument is that even if Searle is right that computers don't think and that, therefore, our brains don't think because they are computers, it remains an open empirical question whether we could build a conscious robot with a digital computer for a brain. And so it remains an open question whether our brains are, in some sense, computers.

Flaws in the Foundations

Marr, Hubel, and Wiesel may have been right that we can gain some insight into how the brain works by thinking of the brain as an information processor. That is, it is reasonable to take up a functional approach to the brain by asking ourselves: Just what sorts of problems is the brain solving? What is it doing? For this reason it is reasonable, as methodology, to approach seeing and other mental powers as information-processing capacities. But these authors seem to have entirely overlooked that this methodological decision on its own did not force them to think of vision as an information-processing problem carried out *in the brain*. Nothing forced the assumption that vision is a process that takes place between the eyeball and the back of the head. And it is that assumption—that the resources for understanding vision, however characterized, are internal to the brain—that dooms the approach. Neural activity simply cannot rise to the level of consciousness, not even when that neural activity is described in computer-theoretic terms.

Hubel and Wiesel's approach to vision is based on the idea that the brain sees by processing special kinds of signals or symbols; the brain sees by building up an internal picture. But the

brain doesn't see; there is no reason to think that seeing happens in the brain. And what good are symbols if no one is around to read them?

Where other scientists were incautious—some trumpeted the existence of "grandmother neurons," i.e., neurons selective for very specific meaningful stimuli like the face of one's grandmother—Hubel and Wiesel were anything but. They hesitated even to refer to their line- and orientation-selective cells as edge or orientation detectors, although that seems to be precisely how they thought of them. Consider this ambivalent passage from Hubel's Nobel lecture:

Orientation-specific simple or complex cells "detect" or are specific for the direction of a short line segment. The cells are thus best not thought of as "line detectors": they are no more line detectors than they are curve detectors. If our perception of a certain line or curve depends on simple or complex cells it presumably depends on a whole set of them, and how the information from such sets of cells is assembled at subsequent stages in the path, to build up what we call "percepts" of lines or curves (if indeed anything like that happens at all), is still a complete mystery.

In this remarkable passage Hubel expresses deep and I think warranted concern about the very theoretical framework that in fact motivates and makes sense of their prizewinning research. If the visual system does not build up a "percept," to use his phrase, on the basis of the sort of processing that their work identifies, then we can be forgiven for asking: Why is it even relevant to an understanding of vision that there are "specialized" cells in the cortex that respond to specific kinds of stimuli of the sort that Hubel and Wiesel discovered? On the assumption that

the visual cortex does build up a representation of the scene on the basis of information in the retina, then the existence of stimulus-selective cells might seem to be tantalizing evidence of how the computational process unfolds. But if it is a "complete mystery" how and whether the brain performs this computational task, as Hubel acknowledges, then it must be admitted that we really have no reason to think Hubel and Wiesel's discoveries tell us anything at all about the brain basis of vision.

This is a harsh conclusion, but one that is hard to avoid. The whole idea that signals are passed from receptors to retinal ganglion cells, up to geniculate cells, and then on to simple, complex and hypercomplex cells, eventually activating the experience of seeing the world, can and should be called into question. As discussed in an earlier chapter, we now know that there are more connections down from "higher" visual areas to lower visual areas than there are connections moving in the opposite direction. That is to say, there is feedback. So whatever is going on, it isn't the sort of simple hierarchical process that Hubel and Wiesel invoke.

We now know that the behavior of cells in the cortex varies, depending on what the animal is doing or what it is paying attention to. The modulation of the behavior of cells depending on the context of the animal's activity is something that Hubel and Wiesel's research did not and could not take into account, for they worked only with animals that were not engaged in any active task: their subjects were unconscious. That is, they were anesthetized, paralyzed, on artificial respiration; stimuli were presented to eyes whose lids were peeled back and held open with clips; eyes were kept moist and clear by means of contact lenses. It is only the assumption that vision is something that happens passively inside the brain that could justify conducting research of this sort in an effort to understand how vision works. But surely we should question this assumption. Remember, we have

no clue how neural activations would or could make visual experience happen. Moreover, it is salutary to remember that animals evolved vision not to represent the world in the head but to enable engaged living—for example, the pursuit of prey and mates and the avoidance of predators and other dangers.

Conclusion: Mind Is Not the Brain's Software

Computers can't think on their own any more than hammers can pound in nails on their own. They are tools we use to think with. For this reason, we make no progress in trying to understand how brains think by supposing that they are computers. In any case, brains don't think: they don't have minds; animals do. To understand the contribution of the brain to the life of the mind, we need to give up once and for all the idea that our minds are achieved inside us by internal goings-on. Once this is clear, we are forced to rethink the value even of Nobel Prize-winning research. This is a disturbing consequence, but one we had better be willing to accept if we want to move forward with a genuinely biological theory of ourselves.