> For your entire life, you've been walking around assuming that your "self" is anchored to a single body that remains stable and permanent at least until death. . . .Yet these [results] suggest the exact opposite—that your body image, despite all its appearance of durability, is an entirely transitory internal construct that can be profoundly altered with just a few simple tricks.[6]

The implications of this discovery for technology-based manipulations of our sense of presence, body, and location are enormous, as we shall see in chapter 4.

To recap, human brains (and indeed those of many other animals) seem to support highly negotiable body-images. As a result, our brains can quite readily project feeling and sensation beyond the biological shell. In much the same way, the blind person's cane or the sports star's racket soon come to feel like genuine extensions of the user's body and senses. Once again, this is because our continual experience of closely correlated action and feedback routines running via these nonbiological peripheries allows the brain to temporarily generate what is really a new kind of "body-image," one that includes the nonbiological components. The transformative effects of this run pretty deep. In a recent neuroscientific experiment in which a monkey repeated food-retrieving actions using a rake, experimenters reported that

> the visual RF [receptive fields] of cells in the anterior bank of the intraparietal sulcus became elongated along an axis of the tool, as if the image of the tool was incorporated into that of the hand.[7]

Otherwise put, the monkey's brain rapidly learned to quite literally treat the rake as an extension of its fingers. It is reasonable to suspect that it is at precisely this point that certain kinds of tools (manually deployed ones) become transparent in use. Here, as elsewhere, the seeds of the most intimate organism-artifact unions are sown by the biological brain itself.

Neural Opportunism

Several other features of our brains combine to make us humans especially open to processes of deep biotechnological symbiosis. One such feature is

what I'll call "neural opportunism." Sit back in your chair and take a look around the room. What did you see? In all likelihood, you had the experience of a succession of rich visual images: images of chairs, books, tables, CDs, audio equipment, whatever. In my own case, I saw a bookshelf stacked rather untidily with too many things I ought to have read, their multicolored spines accusingly flaunting clear, crisp, inviting titles. Looking around, I glimpsed an open closet liberally sprinkled with gaudy Hawaiian shirts, stark against the mundane backdrop of darker, workaday clothing. But now let's ask what turns out to be an especially tricky question. In generating that sequence of visual experiences, what information did my biological brain actually bother to extract and process? The answer is—significantly less than we might have guessed.

To understand this, first reflect that the human visual system supports only a small area of high-resolution processing, corresponding to the fraction of the visual field that falls into central focus. When we inspect a visual scene, our brains actively move this small high-resolution window (the fovea) around the scene, alighting first on one location, then another. The whole of my bookcase, for example, cannot possibly fit into this small foveal area while I remain seated at my desk. My overall visual field (that area *plus* the low-resolution peripheries) is, of course, much larger, and a sizable chunk of my bookshelf falls within this coarse-grained view. It has been known since 1967 that the brain makes very intelligent use of its small high-resolution fovea, moving it around the scene (in a sequence of rapid motions known as visual saccades) in ways delicately suited to the specific problem at hand.[8] This can be seen from the fact that human subjects presented with *identical* pictures, but told to prepare to solve *different* kinds of problems (some might be told to "give the sex and ages of the people in the picture," while others are asked to simply "describe what is going on" and still others to prepare to "recall the objects in the room"), show very different patterns of visual saccade. These saccades, it is also worth commenting, are fast—perhaps three per second—and often repetitive, in that they may visit and revisit the very same part of the scene. What are they for?

One possibility, at this point, is that each saccade is being used to slowly build up a detailed internal representation of the *salient* aspects of the scene. The visual system would thus be selective, but would still be doing what we intuitively expect. It would be using visual input to slowly build up a detailed neural image of the scene. Subsequent research, however, suggests that the real story is even stranger than that. We can get a sense of this even before looking at the scientific experiments, by thinking about some magic tricks.

There is an entertaining web site where you can try out the following trick.[9] You are shown, on screen, a display of six playing cards (new ones are generated each time the trick is run). In the time-honored tradition, you are then asked to mentally select and remember one of those cards. You click on an icon and the cards disappear, to be replaced by a brief "distracter" display. Click again and a five-card (one less) array appears. As if by magic, the very card that you picked is the one that has been removed. How can it be? Could the computer have somehow monitored your eye movements? A version of this trick is displayed on pages 65 and 66 of this book. Go to page 65 and immediately pick a card from the display shown in Fig. 3.1. Concentrate on that card. Remember it. Now go to page 66. Did we remove the very card you chose? Amazing isn't it! I must confess that on first showing (and second, and third) I was quite unable to see how the trick was turned.

Here's the secret. The original array will always show six cards of a similar broad type: six face cards, or six assorted low-ranking cards (between about two and six, for example). When the new, five-card array appears, NONE of these cards will be in the set. But the new five-card array will be of the same type: all face cards, low cards, whatever. In this way, the trick capitalizes on the visual brain's laziness (or efficiency, if you prefer). It seems to the subject exactly as if all that has happened is that one card (the very one he mentally selected!) has disappeared from an otherwise unchanged array. But the impression that the original array is still present is a mistake, rooted in the fact that all the brain had actually encoded was something like "lots of royal cards including my mentally selected king of hearts." Magic tricks such as these rely on our tendency to overestimate what we actually see in a single glance, and on the manipulation of our attention so as to actively inhibit the extraction of crucial information at certain critical moments. The philosopher Daniel Dennett makes a similar point using a different card trick.[10] He invites someone to stand in front of him and fixate on his (Dennett's) nose. In each outstretched arm he holds

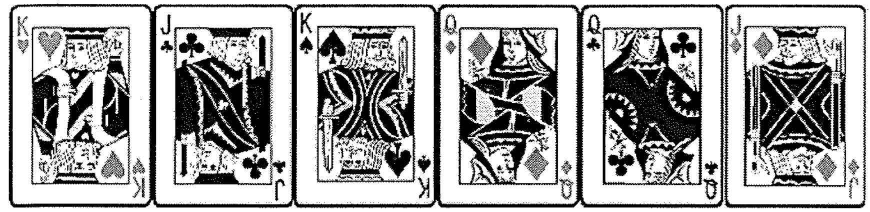

Fig. 3.1 Pick a card and concentrate on it very hard. We will make your card, and only your card, disappear. Turn to page 66 and see if your card is now missing from the array! Thanks to Andy Bauch for permission to show the trick here.

a playing card. He brings his arms in steadily. The question is, at what point will the subject be able to identify the color of the card? Here too, we may be surprised. Color sensitivity, it turns out, is available only in a small and quite central part of the visual field. Yet my conscious experience, clearly, is not of a small central pool of color surrounded by a vague and out-of-focus expanse of halftones. Things look colored all the way out. Once again, it begins to look as if my conscious visual experience is overestimating the amount and quality of information it makes available.

Now imagine that you are the subject of another famous experiment.[11] You are seated in front of a computer screen on which is displayed a page of text. Your eye movements are being automatically tracked and monitored. Your experience, as you report it, is of a solid, stable page of readable text. The experimenter then reveals the trick. In fact, the text to the left and right of a moving "window" has been constantly filled with junk characters, not recognizable English text at all. But because that small window of normal, readable text has been marching in step with your central perceptual span, you never noticed anything odd or unusual. For comparison, this is as if my bookshelf only ever once contained (at the same moment) four or five clearly titled books, and the rest of the titles were all senseless junk. Nonetheless, it would have looked to me as if I were seeing a wide array of clear English titles at all times. In the case of the screen of text, the window of "good stuff" needed to support the illusion is about eighteen characters wide, with the bulk of the characters falling to the right of the point of fixation (probably because English is read left to right).

Similar experiments have been performed using pictures of a visual scene, such as a house, with a parked car and a garden.[12] As before, the victim sits

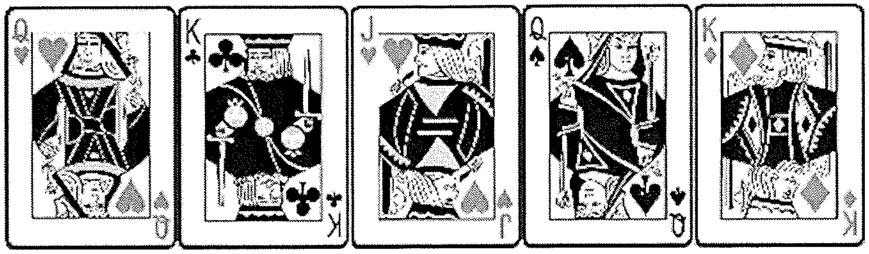

Fig. 3.2 Did we get your card? Puzzled? Go back to page 65 and try again.

in front of a computer-generated display. Her eye movements are monitored and, while they saccade around the display, changes are clandestinely made: the colors of flowers and cars are altered, the structure of the house may be changed; yet these changes, likewise, go undetected. We now begin to understand why the patterns of saccade are not cumulative—why we visit and repeatedly *revisit* the same locations. It is because our brains just don't bother to create rich inner models. Why should they? The world itself is still there, a complex and perfect store of all that data, nicely poised for swift retrieval as and when needed by the simple expedient of visual saccade to a selected location. The kind of knowledge that counts, it begins to seem, is not detailed *knowledge* of what's out there, so much as a broad *idea* of what's out there: one capable of then informing on-the-spot processes of information retrieval and use.

Finally, lest you suspect that these effects (known as "change blindness") are somehow caused by the unnaturalness of the experimental situations, consider some recent work by Dan Simons and Dan Levin.[13] Simons and Levin took this research into the real world. They set up a kind of slapstick scenario in which an experimenter would pretend to be lost on the Cornell campus, and would approach an unsuspecting passerby to ask for directions. Once the passerby started to reply, two people carrying a large door would (rudely!) walk right between the inquirer and the passerby. During the walk through, however, the original inquirer is deftly replaced (under cover of the door) by a different person. Only 50 percent of the subjects (the direction-givers) noticed the change. Yet the two experimenters were of different heights, wore different clothes, had very different voices, and so on. The conclusion that Simons and Levin draw is that our failures to detect change are not due to the artificialness of the computer-

screen experiments. Instead, they arise because "we lack a precise representation of our visual world from one view to the next" and encode only a kind of rough gist of the current scene—enough to support a broad underlying sense of what's going on *insofar as it matters to us*, and enough to guide further intelligent information-retrieval, via directed saccades, as and when needed.[14]

A final demonstration of these startling effects can be obtained using the so-called flicker paradigm. Here, you look at a computer-generated image, which flashes on and off, with a masking screen intervening. Between each showing of the image, something changes. Even when these changes are large and significant (for example, one jet engine of an airplane, shown at center screen, repeatedly appears and then disappears), we do not easily spot them. For many of these changes, subjects need to view the rapidly alternating images for nearly a minute before they see the change. Once they have spotted the change they find it hard to believe that they did not see it at once. Normally, motion cues would alert us to the area of the visual scene where a change was occurring. But in these experiments the motion cue is being screened off by the intervening blank screen (the mask). Without that cue, the changes prove very hard to detect. You can try these experiments out at various sites on the web listed in the note.[15]

What all this suggests is that the visual brain may have hit upon a very potent problem-solving strategy, one that we have already encountered in other areas of human thoughts and reason. It is the strategy of preferring *meta-knowledge* over *baseline knowledge*. Meta-knowledge is knowledge about how to acquire and exploit information, rather than basic knowledge about the world. It is not knowing so much as knowing how to find out. The distinction is real, but the effect is often identical. Having a super-rich, stable inner model of the scene could enable you to answer certain questions rapidly and fluently, but so could knowing how to rapidly retrieve the very same information as soon as the question is posed. The latter route may at times be preferable since it reduces the load on biological memory itself. Moreover, our daily talk and practice often blurs the line between the two, as when we (quite properly) expect others to know what is right in front of their eyes. Or when—to recall an example from the previous chapter—we say that we know the time, before looking, simply because we are wearing a watch!

The visual brain is thus *opportunistic*, always ready to make do and mend, to get the most from what the world already presents rather than building whole inner cognitive routines from neural cloth. Instead of attempting to create, maintain, and update a rich inner representation (inner image or model) of the scene, it deploys a strategy that roboticist Rodney Brooks describes as "letting the world serve as its own best model."[16] Brooks's idea is that instead of tackling the alarmingly difficult problem of using input from a robot's sensors to build up a highly detailed, complex inner model of its local surroundings, a good robot should use sensing frugally in order to select and monitor just a few critical aspects of a situation, relying largely upon the persistent physical surroundings themselves to act as a kind of enduring, external data-store: an external "memory" available for sampling as needs dictate.

Our brains, like those of the mobile robots, try whenever possible to let the world serve as its own best model. In the light of this, some writers have suggested that our daily experience of a rich, highly detailed visual scene unfolding before the mind's eye must be something of an illusion.[17] On this view, it only *seems* to us as if we enjoy rich visual experience, thanks to that rapid capacity to retrieve more detailed information from the world as and when required. I now suspect, however, that this is a rather more delicate call than it at first appears, and the reason is one that bears quite directly on the larger themes of the present treatment.[18]

To see what I mean, let's leave the visual case (temporarily) and consider a very different example. Imagine you are a devout sports fan, and that you know thousands upon thousands of somewhat arcane facts about the performance statistics of players in U.S. women's basketball over the last twenty years. One day, as you are seated on your favorite barstool awaiting the start of a game, conversation turns to the Sacramento Monarchs' Kedra Holland-Corn. You immediately recall a few useful facts: that in 2000, her three-point field goal percentage was .361, ranking her seventeenth in the WNBA; that she scored a staggering twenty-three points in 8-of-12 shooting in a recent win over Los Angeles, and so on. While you reel off these facts and figures, you are implicitly aware that you could have done the same for any number of other players in the WNBA. You are not currently thinking about, for example, Jennifer Azzi of the Utah Starzz. But had the need arisen, her field throw percentage of .930 in the 2000 season

would have been as readily available as the data on Holland-Corn. Hence, we have no hesitation in ascribing to you a rich underlying body of basketball knowledge. It is not that all that knowledge is currently *conscious*. You are not, let us imagine, right now experiencing any thoughts about Jennifer Azzi, only about Kedra Holland-Corn, but you do experience yourself as *in command* of a rich and detailed database in which all that information is stored, organized, and poised for easy recovery and use. Returning to the case of vision, notice that there, too, we find ourselves in command of a rich and detailed visual database in which information about the current scene is stored, organized, and poised for use. It is just that much of the database, in the case of vision, is located *outside* the head and is accessed by outward-looking sensory apparatus, principally the eyes. In each case, however, it is the fact that you can indeed access all this data swiftly and easily as and when required that bears out our judgments about the richness of our own knowledge and understanding.

Word Brains

You can probably see where this is heading and how it fits in with our emerging cyborg theme tune. It just *doesn't matter* whether the data are stored somewhere inside the biological organism or stored in the external world. What matters is how information is poised for retrieval and for immediate use as and when required. Often, of course, information stored outside the skull is not so efficiently poised for access and use as information stored in the head. And often, the biological brain is insufficiently aware of exactly *what* information is stored outside to make maximum use of it; old fashioned encyclopedias suffer from all these defects and several more besides. But the more these drawbacks are overcome, the less it seems to matter (scientifically or philosophically) exactly *where* various processes and data stores are physically located, and whether they are neurally or technologically realized. The opportunistic biological brain doesn't care. Nor—for many purposes—should we.

Consider next the opportunistic infant brain in the ecologically unique environment of spoken and written words. What might the reliable presence of linguistic surroundings do for brains like ours? This is a complex and much-debated issue.[19] But the small thread that I want to pull on here concerns the role of spoken language itself as a kind of triggering cognitive