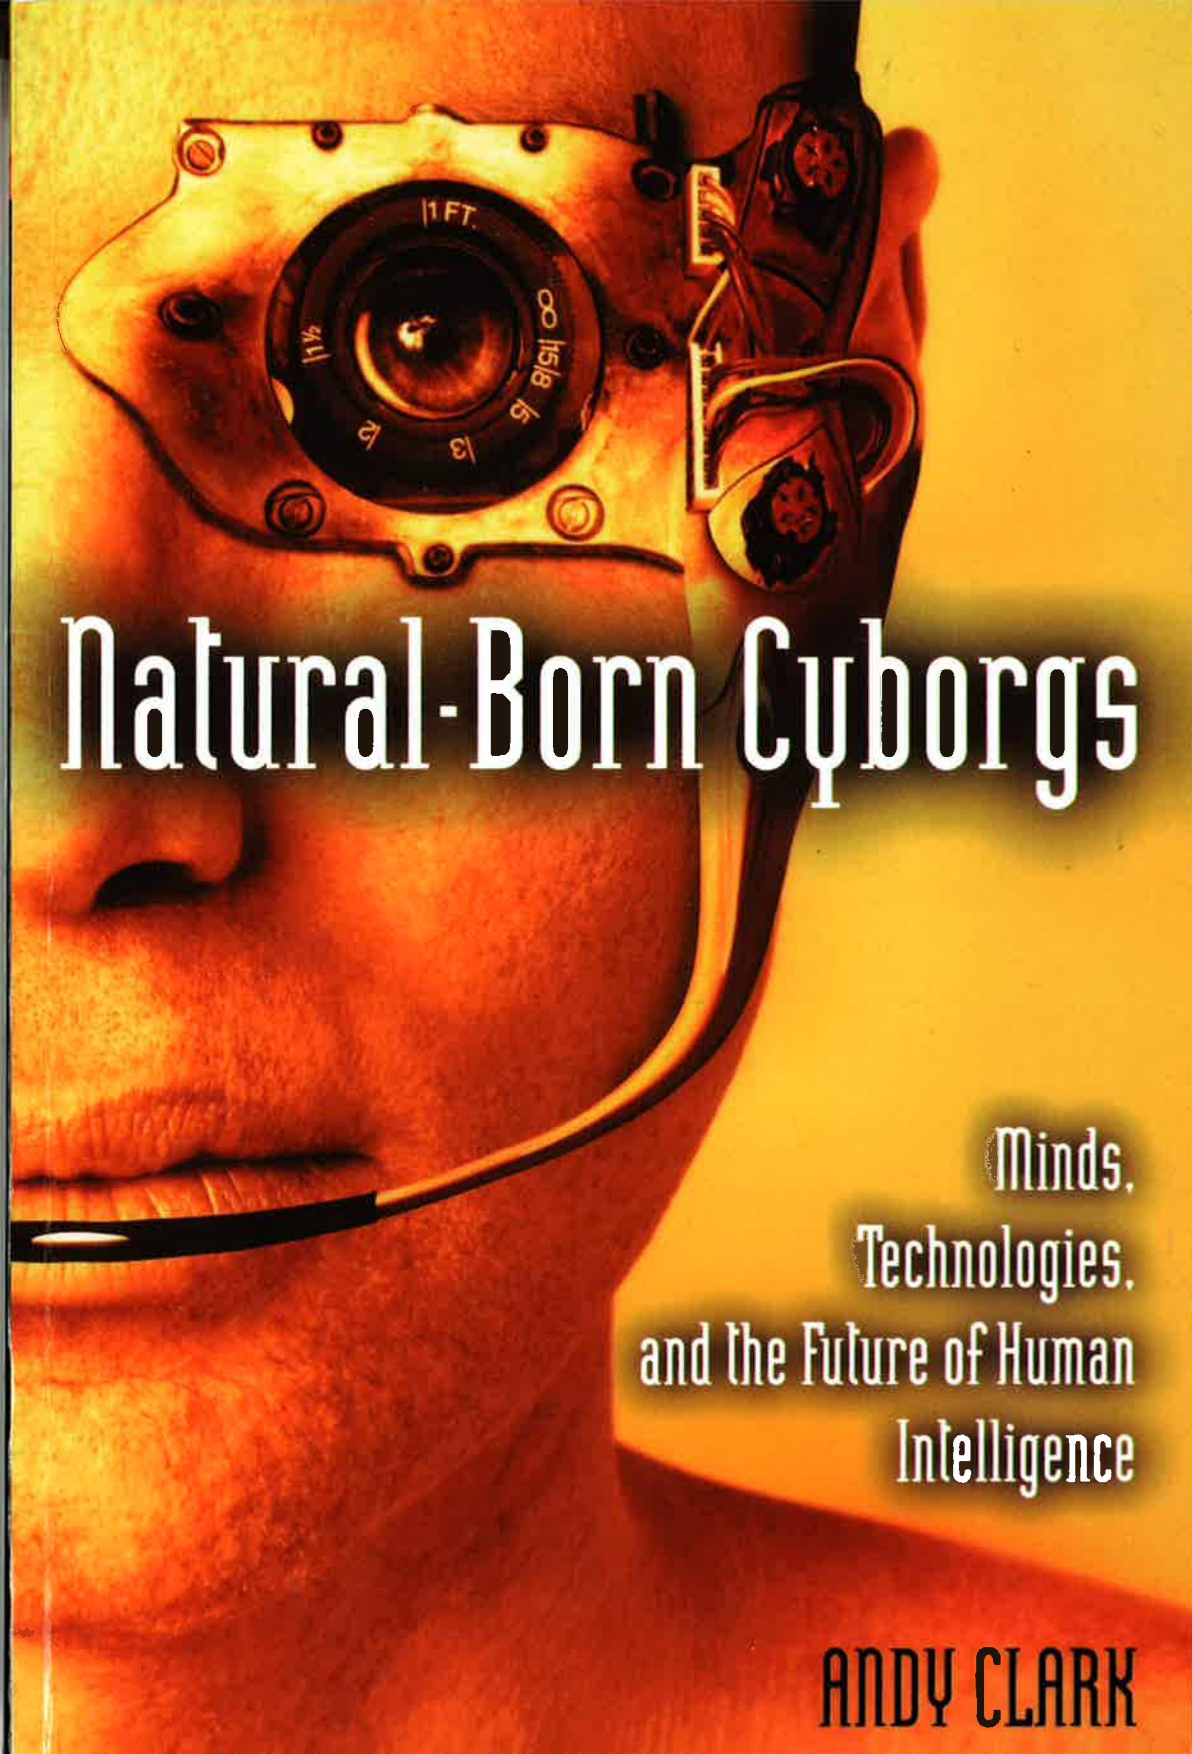

Natural-Born Cyborgs

Minds, Technologies, and the Future of Human Intelligence

ANDY CLARK

The human skin is an artificial boundary: the world wanders into it, and the self wanders out of it, traffic is two-way and constant.

—Bernard Wolfe, *Limbo*

We're here to go.

—William S. Burroughs, *Dead City Radio*

Introduction

The Naked Cyborg

My body is an electronic virgin. I incorporate no silicon chips, no retinal or cochlear implants, no pacemaker. I don't even wear glasses (though I do wear clothes), but I am slowly becoming more and more a cyborg. So are you. Pretty soon, and still without the need for wires, surgery, or bodily alterations, we shall all be kin to the Terminator, to Eve 8, to Cable . . . just fill in your favorite fictional cyborg. Perhaps we already are. For we shall be cyborgs not in the merely superficial sense of combining flesh and wires but in the more profound sense of being human-technology symbionts: thinking and reasoning systems whose minds and selves are spread across biological brain and nonbiological circuitry. This book is the story of that transition and of its roots in some of the most basic and characteristic facts about human nature. For human beings, I want to convince you, are *natural-born* cyborgs.

This may sound like futuristic mumbo-jumbo, and I happily confess that I wrote the preceding paragraph with an eye to catching your attention, even if only by the somewhat dangerous route of courting your immediate disapproval! But I do believe that it is the plain and literal truth. I believe, to be clear, that it is above all a *SCIENTIFIC* truth, a reflection of some deep and important facts about (a whiff of paradox here?) our special, and distinctively *HUMAN*, nature. Certainly I don't think this tendency toward cognitive hybridization is a modern development. Rather, it is an aspect of our humanity, which is as basic and ancient as the use of speech and which has been extending its territory ever since. We see some of the "cognitive fossil trail" of the cyborg trait in the historical procession of potent cognitive technologies that begins with speech and counting, morphs first into written text and numerals, then into early printing (without moveable typefaces), on to the revolutions of moveable typefaces and the printing press, and most recently to the digital encodings that bring text, sound, and image into a uniform and widely transmissible format. Such technologies, once up and running in the various appliances and institutions that surround us, do far more than merely allow for the external storage and transmission of ideas. They constitute, I want to say, a cascade of "mindware upgrades": cognitive upheavals in which the effective architecture of the human mind is altered and transformed.

It was about five years ago that I first realized we were, at least in that specific sense, all cyborgs. At that time I was busy directing a new interdisciplinary program in philosophy, neuroscience, and psychology at Washington University in St. Louis. The realization wasn't painful; it was, oddly, reassuring. A lot of things now seemed to fall into place: why we humans are so deeply different from the other animals, while being, quite demonstrably, not so *very* different in our neural and bodily resources; why it was so hard to build a decent thinking robot; why the recent loss of my laptop had hit me like a sudden and somewhat vicious type of (hopefully transient) brain damage.

I'd encountered the idea that we were all cyborgs once or twice before, but usually in writings on gender or in postmodernist (or post postmodernist) studies of text. What struck me in July 1997 was that this kind of story was the literal and scientific truth. The human mind, if it is to be the physical organ of human reason, simply cannot be seen as bound and restricted by the biological skinbag. In fact, it has *never been* thus restricted and bound, at least not since the first meaningful words were uttered on some ancestral plain. But this ancient seepage has been gathering momentum with the advent of texts, PCs, coevolving software agents, and user-adaptive home and office devices. The mind is just less and less in the head.

If we do not always see this, or if the idea seems outlandish or absurd, that is because we are in the grip of a simple prejudice: the prejudice that

whatever matters about *my* mind must depend solely on what goes on inside my own biological skin-bag, inside the ancient fortress of skin and skull. This fortress has been built to be breached; it is a structure whose virtue lies in part in its capacity to delicately gear its activities in order to collaborate with external, nonbiological sources of order to better solve the problems of survival and reproduction. It is because we are so prone to think that the mental action is all, or nearly all, on the inside, that we have developed sciences and images of the mind that are, in a fundamental sense, inadequate to their self-proclaimed target. So it is actually important to begin to see ourselves aright—it matters for our science, our morals, and our sense of self.

What, then, is the role of the biological brain, of those few pounds of squishy matter in your skull? The squishy matter is great at some things. It is expert at recognizing patterns, at perception, and at controlling physical actions, but it is not so well designed (as we'll see) for complex planning and long, intricate, derivations of consequences. It is, to put it bluntly, bad at logic and good at Frisbee. It is both our triumph and our burden, however, to have created a world so smart that it allows brains like ours to go where no animal brains have gone before. The story I want to tell is the story of that triumph, and of what it means for our understanding of ourselves: dumb thinkers in a smart world, or smart thinkers whose boundaries are simply not those of skin and skull?

The cyborg is a potent cultural icon of the late twentieth century. It conjures images of human-machine hybrids and the physical merging of flesh and electronic circuitry. My goal is to hijack that image and to reshape it, revealing it as a disguised vision of (oddly) our own biological nature. For what is special about human brains, and what best explains the distinctive features of human intelligence, is precisely their ability to enter into deep and complex relationships with nonbiological constructs, props, and aids. This ability, however, does not depend on physical wire-and-implant mergers, so much as on our openness to information-processing mergers. Such mergers may be consummated without the intrusion of silicon and wire into flesh and blood, as anyone who has felt himself thinking *via* the act of writing already knows. The familiar theme of "man the toolmaker" is thus taken one crucial step farther. Many of our tools are not just external props and aids, but they are deep and integral parts of the problem-solving systems we now identify as human intelligence. Such tools are best conceived as proper parts of the computational apparatus that constitutes our minds.

The point is best made by the series of extended concrete examples that I develop in this book. Consider, as a truly simplistic cameo, the process of using pen and paper to multiply large numbers.[1] The brain learns to make the most of its capacity for simple pattern completion (4 × 4 = 16, 2 × 7 = 14, etc.) by acting in concert with pen and paper, storing the intermediate results outside the brain, then repeating the simple pattern completion process until the larger problem is solved. The brain thus dovetails its operation to the external symbolic resource. The reliable presence of such resources may become so deeply factored in that the biological brain alone is rendered unable to do the larger sums.

Some educationalists fear this consequence, but I shall celebrate it as the natural upshot of that which makes us such potent problem-solving systems. It is because our brains, more than those of any other animal on the planet, are primed to seek and consummate such intimate relations with nonbiological resources that we end up as bright and as capable of abstract thought as we are. It is because we are natural-born cyborgs, forever ready to merge our mental activities with the operations of pen, paper, and electronics, that we are able to understand the world as we do. There has been much written about our imminent "post-human" future, but if I am right, this is a dangerous and mistaken image. The very things that sometimes seem most post-human, the deepest and most profound of our potential biotechnological mergers, will reflect nothing so much as their thoroughly human source.

My cat Lolo is not a natural-born cyborg. This is so despite the fact that Lolo (unlike myself) actually does incorporate a small silicon chip. The chip is implanted below the skin of his neck and encodes a unique identifying bar code. The chip can be read by devices common in veterinarians' offices and animal shelters; it identifies me as Lolo's owner so we can be reunited if he is ever lost. The presence of this implanted device makes no difference to the shape of Lolo's mental life or the range of projects and endeavors he undertakes. Lolo currently shows no signs of cat-machine symbiosis, and for that I am grateful. By contrast it is our special character, as human beings, to be forever driven to create, co-opt, annex, and exploit nonbiological props and scaffoldings. We have been designed, by Mother Nature, to exploit deep neural plasticity in order to become one with our

best and most reliable tools. Minds like ours were made for mergers. Tools-R-Us, and always have been.

New waves of user-sensitive technology will bring this age-old process to a climax, as our minds and identities become ever more deeply enmeshed in a nonbiological matrix of machines, tools, props, codes, and semi-intelligent daily objects. We humans have always been adept at dovetailing our minds and skills to the shape of our current tools and aids. But when those tools and aids start dovetailing back—when our technologies actively, automatically, and continually tailor themselves to us just as we do to them—then the line between tool and user becomes flimsy indeed. Such technologies will be less like tools and more like part of the mental apparatus of the person. They will remain tools in only the thin and ultimately paradoxical sense in which my own unconsciously operating neural structures (my hippocampus, my posterior parietal cortex) are tools. I do not really "use" my brain. There is no user quite so ephemeral. Rather, the operation of the brain makes me who and what I am. So too with these new waves of sensitive, interactive technologies. As our worlds become smarter and get to know us better and better, it becomes harder and harder to say where the world stops and the person begins.

Mind-expanding technologies come in a surprising variety of forms. They include the best of our old technologies: pen, paper, the pocket watch, the artist's sketchpad, and the old-time mathematician's slide rule. They include all the potent, portable machinery linking the user to an increasingly responsive world wide web. Very soon, they will include the gradual smartening-up and interconnection of the many everyday objects that populate our homes and offices.

However, this is not primarily a book about new technology. Rather, it is about us, about our sense of self, and about the nature of the human mind. It targets the complex, conflicted, and remarkably ill-understood relationship between biology, nature, culture, and technology. More a work of science-sensitive philosophy than a futurist manifesto, my goal is not to guess at what we might soon become but to better appreciate what we already are: *creatures whose minds are special precisely because they are tailor-made for multiple mergers and coalitions*.

All this adds important complexity to recent evolutionary psychological accounts that emphasize our ancestral environments.[2] We must take very seriously the profound effects of a plastic evolutionary overlay that yields a constantly moving target, an extended cognitive system whose constancy lies mainly in its continual openness to change. Even granting that the biological innovations that got this ball rolling may have consisted only in some small tweaks to an ancestral repertoire, the upshot of this subtle alteration is now a sudden, massive leap in the space of mind design. Our cognitive machinery is now intrinsically geared to self-transformation, artifact-based expansion, and a snowballing/bootstrapping process of computational and representational growth.

The line between biological self and technological world was, in fact, never very firm. Plasticity and multiplicity are our true constants, and new technologies merely dramatize our oldest puzzles (prosthetics and telepresence are just walking sticks and shouting, cyberspace is just one more place to be). Human intellectual history is, in large part, the tale of this fragile and always unstable frontier. The story I tell overlaps some familiar territory, touching on our skills as language-users, toolmakers, and tool-users. But it ends by challenging much of what we think we know about who we are, what we are, and even where we are. It ought to start, perhaps, somewhere on some dusty ancestral savanna, but join me instead on a contemporary city street, abuzz with the insistent trill of a hundred cell phones. . . .

Wired

Brighton main street, hub of a once-sleepy English seaside town lately transformed into a hi-tech haven and club-culture capital. This used to be my town, but it has changed. The shops tell a new story. I walk slowly, taking stock. I count one cell phone shop, one Starbucks, another cell phone shop, a hardware store, *another* cell phone shop, a clothes store, another coffee shop (this one offering full internet access), *yet another* cell phone shop . . .

The toll steadily mounts. Brighton, in my ten-year absence in the United States, has converted itself into a town that seems to sell nothing but coffee and cell phones. The center of town is now home to no fewer than fifty shops dedicated entirely to the selling of cell phones and their contracts. Then there are the various superstores that offer these phones alongside a variety

of other goods. This is quite astonishing. For a relatively small town (around 250,000) this is surely a massive load. Yet business looks good and no wonder: everywhere I turn there are people with phone to ear, or punching in text messages using the fluent two-thumbed touch typing that is the badge of the younger users. Some, with fancier handsets, are using the phone to surf the web. This town is wired.

Not only is it wired. Half the people aren't entirely where they seem to be. I spent last Christmas in the company of a young professional whose phone was hardly ever out of his hands. He wasn't using the phone to speak but was constantly sending or receiving small text messages from his lover. Those thumbs were flying. Here was someone living a divided life: here in the room with us, but with a significant part of him strung out in almost constant, low-bandwidth (but apparently highly satisfying) contact with his distant friend.

The phone of the flying thumbs was a Nokia. Thanks in large part to Nokia (the firm, based in the Finnish town of the same name) the Finns emerged as early heavy-hitters in the European cell phone league. In 1999, 67 percent of the Finnish population owned and used cell phones compared to 28 percent in the United States. And these are not wimpy devices. Nokia is a pioneer of Wireless Application Protocol (WAP) technology, which supports fluent interfacing between the phone and the internet. Top of the line Finnish phones have for many years opened in the middle to reveal a small keyboard and screen supporting full fax, web, and e-mail capability. But it is not the potency of the technology so much as the pregnancy of the slang that really draws me to Finland. Finnish youngsters have dubbed the cell phone "kanny," which means extension of the hand.[3] The mobile is thus both something you use (as you use your hands to write) and something that is part of you. It is like a prosthetic limb over which you wield full and flexible control, and on which you eventually come to automatically rely in formulating and carrying out your daily goals and projects. Just as you take for granted your ability to use your vocal cords to speak to someone in the room beside you, you may take for granted your ability to use your thumbs-plus-mobile to send text to a distant lover. The phone really did seem to be part of the man, and the Finnish slang captures the mood.

I am surprised, but I shouldn't be. As a working cognitive scientist, the more I have learned about the brain and the mind, the more convinced I have become that the everyday notions of "minds" and "persons" pick out deeply plastic, open-ended systems—systems fully capable of including nonbiological props and aids as quite literally parts of themselves. No wonder the cell phone shops were full. These people were not just investing in new toys; they were buying *mindware upgrades*, electronic prostheses capable of extending and transforming their personal reach, thought, and vision.

Upgrades, as we all know, can be mixed blessings. Every new capacity brings new limits and demands. We may, for example, start to spread ourselves too thin, reconfiguring our work and social worlds in new and not necessarily better ways. Certainly, I felt more than a tad jealous of my friend's constant low-bandwidth info-dribble. It took some of him away from those he was physically beside. Later on, we'll take a closer look at some of these pros and cons in our cyborg future.

Brighton main street, then, is just one more sign of the times. As technology becomes portable, pervasive, reliable, flexible, and increasingly personalized, so our tools become more and more a part of who and what we are. With WAP-enhanced cell and access to our own personalized versions of the web in hand we see farther, organize better, know more. The temporary disability caused by a dead battery is unnerving. It seems we just aren't ourselves today. (The loss of my laptop, as I mentioned earlier, underlined this in a painfully personal way. I was left dazed, confused, and visibly enfeebled—the victim of the cyborg equivalent of a mild stroke.) So I, of all people, really *shouldn't* have been surprised. It is our natural proclivity for tool-based extension, and profound and repeated self-transformation, that explains how we humans can be *so very special* while at the same time being not so very different, biologically speaking, from the other animals with whom we share both the planet and most of our genes. What makes us distinctively human is our capacity to continually restructure and rebuild our own mental circuitry, courtesy of an empowering web of culture, education, technology, and artifacts. Minds like ours are complex, messy, contested, permeable, and constantly up for grabs. The neural difference that makes all this possible is probably not very large, but its effects are beyond measure.

Don't believe it yet? Or don't think it matters anyway? Both are fair and proper responses. I began deliberately with a technology—the cell phone—which is at once familiar yet insufficiently fluid and user-responsive to make

(as yet) the strongest possible kind of case. And I have rehearsed none of the interlocking evidence (some philosophical, some psychological, some neuroscientific), which actually led me to embrace such a strong thesis in the first place.

Before the day is done, however, I hope to convince you at least of this: that the old puzzle, the mind-body problem, really involves a hidden third party. It is the mind-body-*scaffolding* problem. It is the problem of understanding how human thought and reason is born out of looping interactions between material brains, material bodies, and complex cultural and technological environments. We create these supportive environments, but they create us too. We exist, as the thinking things we are, only thanks to a baffling dance of brains, bodies, and cultural and technological scaffolding. Understanding this evolutionarily novel arrangement is crucial for our science, our morals, and our self-image both as persons and as a species.

CHAPTER 1

Cyborgs Unplugged

Rats in Space

The year is 1960. The pulse of space travel beats insistently within the temples of research and power, and the journal *Astronautics* publishes the paper that gave the term "cyborg" to the world.[1] The paper, titled "Cyborgs and Space," was based on a talk, "Drugs, Space and Cybernetics," presented that May to the Air Force School of Aviation Medicine in San Antonio, Texas. The authors were Manfred Clynes and Nathan Kline, both working for the Dynamic Simulation Laboratory (of which Kline was director) at Rockland State Hospital, New York. What Clynes and Kline proposed was simply a nice piece of lateral thinking. Instead of trying to provide artificial, earth-like environments for the human exploration of space, why not alter the humans so as to better cope with the new and alien demands? "Space travel," the authors wrote, "challenges mankind not only technologically, but also spiritually, in that it invites man to take an active part in his own biological evolution."[2] Why not, in short, reengineer the humans to fit the stars?

In 1960, of course, genetic engineering was just a gleam in science fiction's prescient eye. And these authors were not dreamers, just creative scientists engaged in matters of national (and international) importance. They were scientists, moreover, working and thinking on the crest of two major waves of innovative research: work in computing and electronic data-processing,[3] and work on cybernetics[4]—the science of control and communication in animals and machines. The way to go, they suggested, was to combine cybernetic and computational approaches so as to create man-machine hybrids, "artifact-organism systems" in which implanted electronic devices use bodily feedback signals to automatically regulate wakefulness, metabolism, respiration, heart rate, and other physiological functions in ways suited to some alien environment. The paper discussed specific artificial interventions that might enable a human body to bypass lung-based breathing, to compensate for the disorientations caused by weightlessness, to alter heart rate and temperature, reduce metabolism and required food intake, and so on.

It was Manfred Clynes who actually first suggested the term "cyborg." Clynes was at that time chief research scientist at Rockland State Hospital and an expert on the design and development of physiological measuring equipment. He had already received a prestigious Baker Award for work on the control of heart rate through breathing and would later invent the CAT computer, which is still used in many hospitals today. When Clynes coined the term "cyborg" to describe the kind of hybrid artifact-organism system they were envisaging, Kline remarked that it sounded "like a town in Denmark."[5] But the term was duly minted, and the languages of fact and fiction permanently altered. Here is the passage as it appeared in *Astronautics*:

> For the exogenously extended organizational complex . . . we propose the term "cyborg." The Cyborg deliberately incorporates exogenous components extending the self-regulating control function of the organism in order to adapt it to new environments.[6]

Thus, amid a welter of convoluted prose, was born the cyborg. The acronym "cyborg" stood for Cybernetic Organism or Cybernetically Controlled Organism; it was a term of art meant to capture both a notion of human-machine merging and the rather specific nature of the merging envisaged. Cyberneticists were especially interested in "self-regulating systems." These are systems in which the results of the system's own activity are "fed back" so as to increase, stop, start, or reduce the activity as conditions dictate. The flush/refill mechanism of a standard toilet is a homey example, as is the thermostat on the domestic furnace. The temperature drops, a circuit is activated, and the furnace comes to life. The temperature rises, a circuit is broken, and the furnace ceases to operate. Even more prosaically, the

toilet is flushed, the ballcock drops, which causes the connected inlet valve to open. Water then flows in until the ballcock, riding on the rising tide, reaches a preset level and thus recloses the valve. Such systems are said to be homeostatically controlled because they respond automatically to deviations from a baseline (the norm, stasis, equilibrium) in ways that drag them back toward that original setting—the full cistern, the preset ambient temperature, and the like.

The human autonomic nervous system, it should be clear, is just such a self-regulating homeostatic engine. It works continuously, and without conscious effort on our part, in order to keep key physiological parameters within certain target zones. As effort increases and blood oxygenation falls, we breathe harder and our hearts beat faster, pumping more oxygen into the bloodstream. As effort decreases and blood oxygen levels rise, breathing and heart rate damp down, reducing the intake and uptake of oxygen.

With all this in mind, it is time to meet the first duly-accredited-and-labeled cyborg. Not a fictional monster, not even a human being fitted with a pacemaker (although they are cyborgs of this simple stripe too), but a white laboratory rat trailing an ungainly appendage—an implanted Rose osmotic pump. This rat (see fig 1.1) was introduced in the 1960 paper by Clynes and Kline as "one of the first cyborgs" and the snapshot, as Donna Haraway wonderfully commented "belongs in Man's family album."[7]

Sadly, the rat has no name, but the osmotic pump does. It is named after its inventor, Dr. Rose, who recently died after a very creative life devoted to

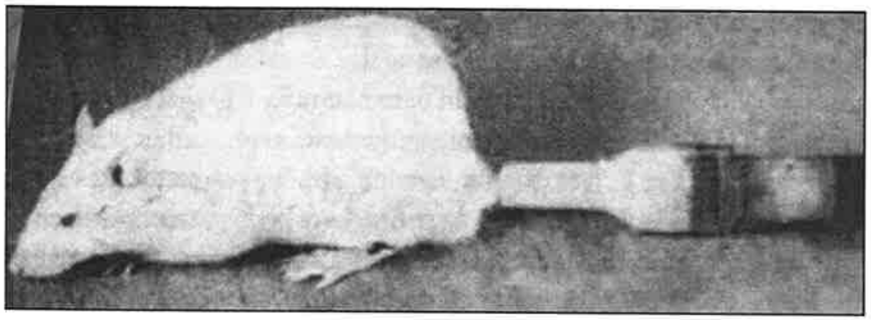

Fig. 1.1 An early (ca. 1955) classic cyborg: rat with implanted Rose osmotic pump. The pump automatically injects chemicals into the rat to form a biotechnological control loop, which can be adapted to unusual conditions (for example, survival in space). By kind permission of Manfred Clynes.

the search for a cure for cancer. So let's respectfully borrow that, calling the whole rat-pump system Rose. Rose incorporates a pressure pump capsule capable of delivering injections at a controlled rate. The idea was to combine the implanted pump with an artificial control loop, creating in Rose a new layer of homeostasis. The new layer would operate like the biological ones without the need for any conscious attention or effort and might be geared to help Rose deal with specific extraterrestrial conditions. The authors speculate, for example, that the automatic, computerized control loop might monitor systolic blood pressure, compare it to some locally appropriate reference value, and administer adrenergic or vasodilatory drugs accordingly.

As cyborgs go, Rose, like the human being with the pacemaker, is probably a bit of a disappointment. To be sure, each incorporates an extra artificial layer of unconsciously regulated homeostatic control. But Rose remains pretty much a rat nonetheless, and one pacemaker doth not a Terminator make. Cyborgs, it seems, remain largely the stuff of science fiction, forty-some years of research and development notwithstanding.

Implant & Mergers

Or do they? Consider next the humble cochlear implant. Cochlear implants, which are already widely in use, electronically stimulate the auditory nerve. Such devices enable many profoundly deaf humans to hear again. However, they are currently limited by requiring the presence of a healthy, undegenerated auditory nerve. A Pasadena-based research group led by Douglas McCreery of Huntington Medical Research Institutes recently addressed this problem by building a new kind of implant (fig 1.2) that bypasses the auditory nerve and connects directly to the brain stem. Earlier versions of such devices have, in fact, been in use for a while, but performance was uninspiring. Uninspiring because these first wave brain stem implants used only an array of surface contacts—flat electrodes laid upon the surface of the brain stem near the ventral cochlear nucleus. The auditory discrimination of frequencies, however, is mediated by stacked layers of neural tissue within the nucleus. To utilize frequency information (to discriminate pitch) you need to feed information differentially into the various layers of this neural structure, where the stimulation of deeper layers results in the audi-

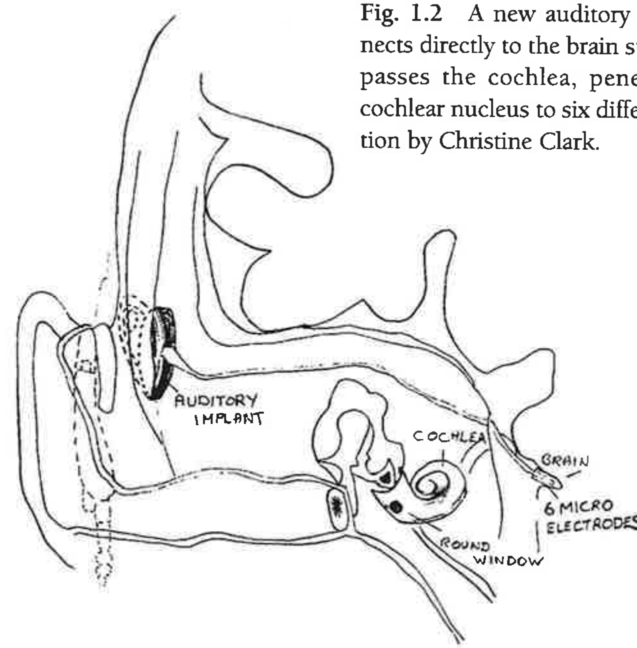

Fig. 1.2 A new auditory prosthesis that connects directly to the brain stem. The implant bypasses the cochlea, penetrating the ventral cochlear nucleus to six different depths. Illustration by Christine Clark.

tory perception of higher frequencies, and so on. The implant being pioneered by McCreery thus reaches deeper than those older, surface contact models, terminating in six iridium microelectrodes each of which penetrates the brain stem to a different depth. The overall system comprises an external speech processor with a receiver implanted under the scalp, directly wired to six different depths within the ventral cochlear nucleus. A Huntington Institute cat, according to neuroscientist and science writer Simon LeVay,[8] is already fitted with the new system and thus joins Rose in our Cyborg Hall of Fame.

The roll call would not be complete, however, without a certain maverick professor. Our next stop is thus the Department of Cybernetics at the University of Reading, in England. It is somewhat of a surprise to find, nowadays, a department of Cybernetics at all. They mostly died out in the early 1960s, to be replaced by departments of Computer Science, Cognitive Science, and Artificial Intelligence. But the real surprise is to find, within this Department of Cybernetics, a professor determined to turn himself into a good old-fashioned flesh-and-wires cyborg. The professor's name is Kevin Warwick, and in his own words:

I was born human. But this was an accident of fate—a condition merely of time and place. I believe it's something we have the power to change.[9]

Warwick began his personal transformation back in 1998, with the implantation of a fairly simple silicon chip, encased in a glass tube, under the skin and on top of the muscle in his left arm. This implant sent radio signals, via antennae placed strategically around the department, to a central computer that responded by opening doors as he approached, turning lights on and off, and so on. This was, of course, all pretty simple stuff and could have been much more easily achieved by the use of a simple device (a smart-badge or card) strapped to his belt or pinned to his lapel. The point of the experiment, however, was to test the capacity to send and receive signals via such an implant. It worked well, and Warwick reported that even in this simple case he quickly came to feel "like the implant was one with my body," to feel, indeed, that his biological body was just one aspect of a larger, more powerful and harmoniously operating system. He reported that it was hard to let go of the implant when the time came for its removal.

The real experiment took place on March 14, 2002, at 8:30 in the morning at the Radcliffe Infirmary, Oxford. There, Warwick received a new and more interesting implant. This consisted of a 100-spike array (see fig. 1.3).

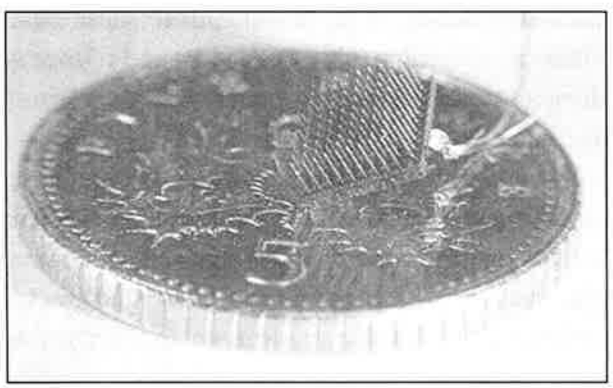

Fig. 1.3 One-hundred-spike array, implanted into Professor Kevin Warwick, March 14, 2002 (shown against a small coin). By kind permission of Professor Warwick and of icube.co.uk.

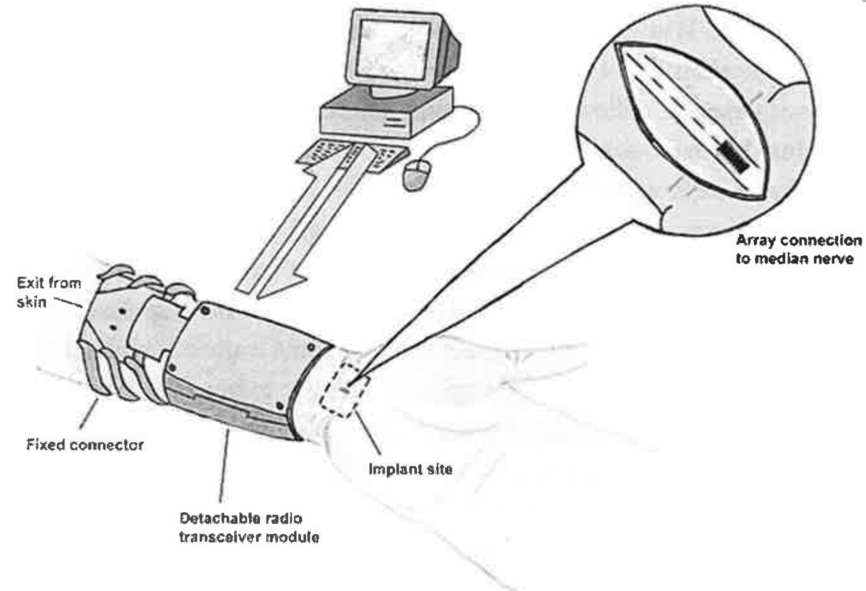

Fig. 1.4 Diagram of implant used by Professor Kevin Warwick. By kind permission of Professor Warwick.

Each of the 100 tips in the array makes direct contact with nerve fibers in the wrist and is linked to wires that tunnel up Professor Warwick's arm, emerging through a skin puncture where they are linked to a radio transmitter/receiver device (fig. 1.4). This allows the median nerve in the arm to be linked by radio contact to a computer. The nerve impulses running between brain and hand can thus be "wiretapped" and the signals copied to the computer. The process also runs in the other direction, allowing the computer to send signals (copies or transforms of the originals) to the implant, which in turn feeds them into the nerve bundles running between Warwick's hand and brain.

The choice of nerve bundles in the arm as interface point is doubtless a compromise. The surgical risks of direct neural interfacing are still quite high (the kind of brain stem implant described earlier, for example, is performed only on patients already requiring surgery to treat neurofibromatosis type 2). But the nerve bundles running through the arm do carry tremendous quantities of information to and from the brain, and they are implicated not just in reaching and grasping but also in the neurophysiology of pain, pleasure, and emotion. Warwick has embarked upon a staged sequence of experiments, the simplest of which is to record and identify the signals associated with specific willed hand motions. These signals can then be played back into his nervous system later on. Will his hand then move again? Will he feel as if he is willing it to move?

The experiment can be repeated with signals wiretapped during episodes of pain or pleasure. Warwick himself is fascinated by the transformative potential of the technology and wonders whether his nervous system, fed with computer-generated signals tracking some humanly undetectable quantity, such as infrared wavelengths, could learn to perceive them, yielding some sensation of seeing or feeling infrared (or ultraviolet, or x-rays, or ultrasound).[10]

Recalling the work on deep (cochlear nucleus penetrating) auditory repair, this kind of thing begins to seem distinctly feasible. Imagine, for example, being fitted with artificial sensors, tuned to detect frequencies currently beyond our reach, but sending signals deep into the developing ventral cochlear nucleus. Human neural plasticity, as we'll later see, may well prove great enough to allow our brains to learn to make use of such new kinds of sensory signal. Warwick is certainly enthusiastic. In his own words, "few people have even had their nervous systems linked to a computer, so the concept of sensing the world around us using more than our natural abilities is still science fiction. I'm hoping to change that."[11]

Finally, in a dramatic but perhaps inevitable twist, there is a plan (if all goes well) to subsequently have a matching but surface-level device connected to his wife, Irena. The signals accompanying actions, pains, and pleasures could then be copied between the two implants, allowing Irena's nervous system to be stimulated by Kevin's and vice versa. The couple also plans to try sending these signals over the internet, perhaps with one partner in London while the other is in New York.

None of this is really science fiction. Indeed, as Warwick is the first to point out, a great deal of closely related work has already been done. Scientists at the University of Tokyo have been able to control the movements of a live cockroach by hooking its motor neurons to a microprocessor; electronically mediated control of some muscular function (lost due to damage or disease) has been demonstrated in several laboratories; a paralyzed stroke patient, fitted with a neurally implanted transmitter, has been able to will a

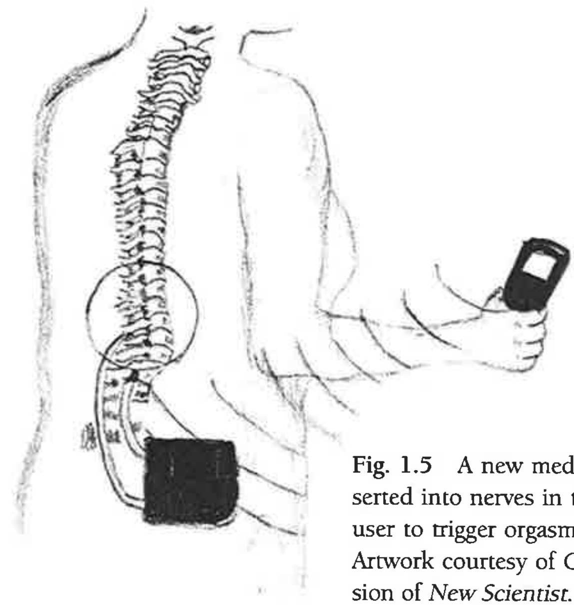

Fig. 1.5 A new medical implant, surgically inserted into nerves in the spinal cord, allows the user to trigger orgasms at the push of a button. Artwork courtesy of Christine Clark, by permission of *New Scientist.*

cursor to move across a computer screen; and rats with similar implants have learned to depress a reward-generating lever by just thinking about it.[12] There is even (fig. 1.5) a female orgasm-generating electronic implant (controlled by a hand-held remote) involving contacts surgically inserted into specific nerves in the spinal cord.[13] Without much doubt, direct bioelectronic signal exchanges, made possible by various kinds of implant technology, will soon open up new realms of human-computer interaction and facilitate new kinds of human-machine mergers. These technologies, for both moral and practical reasons, will probably remain, in the near future, largely in the province of restorative medicine or military applications (such as the McDonnell-Douglas Advanced Tactical Aircraft Program, which envisages a fighter plane pilot whose neural functions are linked directly into the on-board computer).[14]

Despite this, genuinely cyborg technology is all around us and is becoming more and more a part of us every day. To see why, we must reflect some more on what really matters even about the classic (wire-and-implant-dominated) cyborg technologies just reviewed. These classic cases all display direct (wire-based) animal-machine interfacing. Much of the thrill, or horror, depends on imagining all those wires, chips, and transmitters grafted onto pulsing organic matter. But what we should really *care about* is not the mere fact of deep implantation or flesh-to-wire grafting, but the complex and transformative nature of the animal-machine relationships that may or may not ensue. And once we see *that*, we open our eyes to a whole new world of cyborg technology.

Recall the case of the cochlear implants, and notice now the particular shape of this technological trajectory. It begins with simple cochlear implants connected to the auditory nerve—just one step up, really, from hearing aids and ear trumpets. Next, the auditory nerve is bypassed, and signals fed to contacts on the surface of the brain stem itself. Then, finally—classic cyborg heaven—microelectrodes actually penetrate the ventral cochlear nucleus itself at varying depths. Or consider Professor Warwick, whose first implant struck us as little more than a smart badge, worn inside the arm. My sense is that as the bioelectronic interface grows in complexity and moves inward, deeper into the brain and farther from the periphery of skin, bone, and sense organs, we become correlatively less and less resistant to the idea that we are trading in genuine cyborg technology.

But just why do we feel that depth *matters* here? It is, after all, pretty obvious that the physical depth of an implant, in and of itself, is insignificant. Recall my microchipped cat, Lolo. Lolo is, by all accounts, a disappointing cyborg. He incorporates a nonbiological component, conveniently placed within the relatively tamper-proof confines of the biological skin (and fur) bag. But he seems determinedly nontransformed by this uninvited bar coding. He is far from anyone's ideal of the cyborg cat. It would make no difference to *this* intuition, surely, were we to implant the bar code chip as deeply as we like—perhaps right in the center of his brain—humane technology and better bar code readers permitting. What we care about, then, is not depth of implanting per se. Instead, what matters to us is the nature and transformative potential of the bioelectronic coalition that results.

Still, the idea that truly profound biotechnological mergers must be consummated deep within the ancient skin-bag runs deep. It is the point source of the undeniable gut appeal of most classic cyborg technologies, whether real or imaginary. Think of adamantium skeletons, skull-guns, cochlear implants, retinal implants, human brains directly "jacked in" to the matrix of cyberspace—the list goes on and on.[15] The deeper within the biological

skin-bag the bioelectronic interface lies, the happier we are, it seems, to admit that we confront a genuine instance of cyborg technology.

Intuitions, however, are strange and unstable things. Take the futuristic topless dancer depicted in Warren Ellis's wonderful and extraordinary *Transmetropolitan*.[16] The dancer (fig. 1.6) displays a fully functional three-inch-high bar code tattooed across both breasts. In some strange way, this merely superficially bar-coded dancer strikes me as a more unnerving, more genuinely cyborg image, than does the bar-coded cat. And this despite the fact that it is the latter who incorporates a genuine "within the skin-bag" implant. The reason for this reaction, I think, is that the image of the bar-coded topless dancer immediately conjures a powerful (and perhaps distressing) sense of a deeply transformed kind of human existence. The image foregrounds our potential status as trackable, commercially interesting sexual units, subject to repeated and perhaps uninvited electronic scrutiny. We resonate with terror, excitement, or both to the idea of ever-deeper neural and bodily implants in part *because* we sense some rough-and-ready (not foolproof, more of a rule-of-thumb) correlation between depth-of-interface and such transformative potential. The deep ventral cochlear nucleus penetrating implants can, after all, upgrade the functionality of certain profoundly deaf patients in a much more dramatic, reliable, and effective fashion than its predecessors. What really counts is a kind of double whammy implicit in the classic cyborg image. First, we care about the potential of technology to become integrated so deeply and fluidly with our existing biological capacities and characteristics that we feel no boundary between ourselves and the

Fig. 1.6 Bar-coded dancer by Warren Ellis and Darick Robertson (detail from *Transmetropolitan 3*, Helix, DC Comics). By kind permission of Warren Ellis and Darick Robertson.

nonbiological elements. Second, we care about the potential of such human-machine symbiosis to transform (for better or for worse) our lives, projects, and capacities.

A symbiotic relationship is an association of mutual benefit between different kinds of entities, such as fungi and trees. Such relationships can become so close and important that we tend to think of the result as a single entity. Lichen, for example, are really symbiotic associations between an alga and a fungus. It is often a vexed question how best to think of specific cases.[17] The case of cognitive systems is especially challenging since the requirement—(intuitive enough for noncognitive cases)—of physical cohesion within a clear inner/outer boundary seems less compelling when information flows (rather than the flow of blood or nutrients) are the key concern.

The traditional twin factors (of contained integration and profound transformation) come together perfectly in the classic cyborg image of the human body deeply penetrated by sensitively interfaced and capacity-enhancing electronics. But in the cognitive case, it is worth considering that what really matters might be just the *fluidity* of the human-machine integration and the resulting *transformation* of our capacities, projects, and lifestyles. It is then an empirical question whether the greatest usable bandwidth and potential lies with full implant technologies or with well-designed nonpenetrative modes of personal augmentation.[18] With regard to the critical features just mentioned, I believe that the most potent near-future technologies will be those that offer integration and transformation *without* implants or surgery: human-machine mergers that simply bypass, rather than penetrate, the old biological borders of skin and skull.

To see what I mean, let us return to the realms of the concrete and the everyday, scene-shifting to the flight deck of a modern aircraft. The modern flight deck, as the cognitive anthropologist Ed Hutchins has pointed out,[19] is designed as a single extended system made up of pilots, automated "fly-by-wire" computer control systems, and various high-level loops in which pilots monitor the computer while the computer monitors the pilots. The shape of these loops is still very much up for grabs. In the European Airbus,[20] the computer pretty much has the final say. The pilot moves the control stick, but the onboard electronics keep the flight deviations inside a preset envelope. The plane is not allowed, no matter what the pilots do with the control stick, to bank more than 67 degrees or to

point the nose upward at more than 30 degrees. These computer-controlled limits are meant to keep the pilots' maneuvers from compromising the planes' structural integrity or initiating a stall. In the Boeing 747-400,[21] by contrast, the pilots still have the final say. In each case, however, under normal operating conditions, large amounts of responsibility are devolved to the computer-controlled autosystem. (The high-technology theorist and science writer Kevin Kelly nicely notes that human pilots are increasingly referred to, in professional training and talk, as "system managers.")[22]

Piloting a modern commercial airliner, it seems clear, is a task in which human brains and bodies act as elements in a larger, fluidly integrated, biotechnological problem-solving matrix. But still, you may say, this is state-of-the-art high technology. Perhaps there is a sense in which, at least while flying the plane, the pilots participate in a (temporary) kind of cyborg existence, allowing automated electronic circuits to, in the words of Clynes and Kline "provide an organizational system in which [certain] problems are taken care of automatically."[23] But most of us don't fly commercial airliners and are not even cyborgs for a day.

A Day in the Life

Or are we? Let's shift the scene again, this time to your morning commute to the office. At 7:30 A.M. you are awoken not by your native biorhythms but by your preset electronic alarm clock. By 8:30 A.M. you are on the road. It is a chilly day and you feel the car begin to skid on a patch of ice. Luckily, you have traction control and the Automatic Braking System (ABS). You simply hit the brakes, and the car takes care of most of the delicate work required. In fact, as we'll see in later chapters, the human brain is a past master at devolving responsibility in just this kind of way. You may consciously decide, for example, to reach for the wine glass. But all the delicate work of generating a sequence of muscle commands enabling precise and appropriate finger motions and gripping is then turned over to a dedicated, unconscious subsystem—a kind of on-board servomechanism not unlike those ABS brakes.

Arriving at your office, you resume work on the presentation you were preparing for today's meeting. First, you consult the fat file of papers marked "Designs for Living." It includes your own previous drafts, and a lot of work by others, all of it covered in marginalia. As you reinspect (for the umpteenth time) this nonbiological information store, your onboard wetware (i.e., your brain) kicks in with a few new ideas and comments, which you now add as supermarginalia on top of all the rest. Repressing a sigh you switch on your Mac G4, once again exposing your brain to stored material and coaxing it, once more, to respond with a few fragmentary hints and suggestions. Tired already—and it is only 10 A.M.—you fetch a strong espresso and go about your task with renewed vigor. You now position your biological brain to respond (piecemeal as ever) to a summarized list of key points culled from all those files. Satisfied with your work you address the meeting, presenting the final plan of action for which (you believe, card-carrying materialist that you are) your biological brain must be responsible. But in fact, and in the most natural way imaginable, your naked biological brain was no more responsible for that final plan of action than it was for avoiding the earlier skid. In each case, the real problem-solving engine was the larger, biotechnological matrix comprising (in the case at hand) the brain, the stacked papers, the previous marginalia, the electronic files, the operations of search provided by the Mac software, and so on, and so on. What the human brain is best at is learning to be a team player in a problem-solving field populated by an incredible variety of nonbiological props, scaffoldings, instruments, and resources. In this way ours are *essentially* the brains of natural-born cyborgs, ever-eager to dovetail their activity to the increasingly complex technological envelopes in which they develop, mature, and operate.

What blinds us to our own increasingly cyborg nature is an ancient western prejudice—the tendency to think of the mind as so deeply special as to be distinct from the rest of the natural order. In these more materialist times, this prejudice does not always take the form of belief in soul or spirit. It emerges instead as the belief that there is something absolutely special about the cognitive machinery that happens to be housed within the primitive bioinsulation (nature's own duct-tape!) of skin and skull. What goes on in there is so special, we tend to think, that the only way to achieve a true human-machine merger is to consummate it with some brute-physical interfacing performed behind the bedroom doors of skin and skull.

However, there is nothing quite *that special* inside. The brain is, to be sure, an especially dense, complex, and important piece of cognitive machinery. It is in many ways special, but it is not special in the sense of providing a privileged arena such that certain operations must occur *inside* that arena, or in directly wired contact with it, on pain of not counting as part of our mental machinery at all. We are, in short, in the grip of a seductive but quite untenable illusion: the illusion that the mechanisms of mind and self can ultimately unfold only on some privileged stage marked out by the good old-fashioned skin-bag. My goal is to dispel this illusion, and to show how a complex matrix of brain, body, and technology can actually constitute the problem-solving machine that we should properly identify as *ourselves*. Seen in this light, the cell phones of the Introduction were not such a capricious choice of entry-point after all. None of us, to be sure, are yet likely to *think* of ourselves as born-again cyborgs, even if we invest in the most potent phone on the market and integrate its sweeping functionality deep into our lives. But the cell phone is, indeed, a prime, if entry-level, cyborg technology. It is a technology that may, indeed, turn out to mark a crucial transition point between the first (pen, paper, diagrams, and digital media dominated) and the second waves (marked by more personalized, online, dynamic biotechnological unions) of natural-born cyborgs.

Already, plans are afoot to use our cell phones to monitor vital signs (breathing and heart rate) by monitoring the subtle bounceback of the constantly emitted microwaves off of heart and lungs.[24] There is a simpler system, developed by the German company Biotronic, and already under trial in England, that uses an implanted sensor in the chest to monitor heart rate, communicating data to the patient's cell phone. The phone then automatically calls for help if heart troubles are detected. The list goes on.[25] The very designation of the mobile unit as primarily a phone is now in doubt, as more and more manufacturers see it instead as a multifunctional electronic bridge between the bearer and an invisible but potent universe of information, control, and response. At the time of writing, the Nokia 5510 combines phone, MP3 music player, FM radio, messaging machine, and game console, while Handspring's Trio incorporates a personal digital assistant. Sony Ericsson's T68i has a digital camera allowing the user to transmit or store color photos. Cell phones with integrated Bluetooth wireless technology (or similar) microchips will be able to exchange information automatically with nearby Bluetooth-enabled appliances. So enabled, a quick call home will allow the home computer to turn on or off lights, ovens, and other appliances.[26] In many parts of the world, the cell phone is already as integral to the daily routines of millions as the wristwatch—that little invention that let individuals take real control of their daily schedule, and without which many now feel lost and disoriented. And all this (in most cases) without a single incision or surgical implant. Perhaps, then, it is only our metabolically based obsession with our own skin-bags that has warped the popular image of the cyborg into that of a heavily electronically penetrated human body: a body dramatically transformed by prostheses, by neural implants, enhanced perceptual systems, and the full line of Terminator fashion accessories. The mistake—and it is a familiar one—was to assume that the most profound mergers and intimacies must always involve literal penetrations of the skin-bag.

Dovetailing

Nonpenetrative cyborg technology is all around us and is poised on the very brink of a revolution. By nonpenetrative cyborg technology I mean all the technological tricks and electronic aids that, as hinted earlier, are already transforming our lives, our projects, and our sense of our own capacities. What mattered *most*, even where dealing with real bioelectronic implants, was the potential for fluid integration and personal transformation. And while direct bioelectronic interfaces may contribute on both scores, there is another, equally compelling and less invasive, route to successful human-machine merger. It is a route upon which we as a society have already embarked, and there is no turning back. Its early manifestations are already part of our daily lives, and its ultimate transformative power is as great as that of its only serious technological predecessor—the printed word. It is closely related to what Mark Weiser, working at XeroxPARC back in 1988, first dubbed "ubiquitous computing" and what Apple's Alan Kay terms "Third Paradigm" computing.[27] More generally, it falls under the category of transparent technologies. Transparent technologies are those tools that become so well fitted to, and integrated with, our own lives and projects that they are (as Don Norman,[28] Weiser, and others insist) pretty much invisible-in-use. These tools or resources are usually no more the object of our conscious thought and reason than is the pen with which we write, the hand that holds it while writing, or the various neural subsystems

that form the grip and guide the fingers. All three items, the pen, the hand, and the unconsciously operating neural mechanisms, are pretty much on a par. And it is this parity that ultimately blurs the line between the intelligent system and its best tools for thought and action. Just as drawing a firm line in this sand is unhelpful and misguided when dealing with our basic biological equipment so it is unhelpful and misguided when dealing with transparent technologies. For instance, do I merely *use* my hands, my hippocampus, my ventral cochlear nucleus, or are they part of the system—the "me"—that does the using?) There is no merger so intimate as that which is barely noticed.

Weiser's vision, ca. 1991, of ubiquitous computing was a vision in which our home and office environments become progressively more intelligent, courtesy of multiple modestly powerful but amazingly prolific intercommunicating electronic devices. These devices, many of which have since been produced and tested at XeroxPARC and elsewhere, range from tiny tabs to medium size pads to full size boards. The tabs themselves will give you the flavor. The idea of a tab is to "animate objects previously inert." Each book on your bookshelf, courtesy of its continuously active tab, would know where it is by communicating with sensors and transmitting devices in the building and office, what it is about, and maybe even who has recently been using it. Anyone needing the book can simply poll it for its current location and status (in use or not). It might even emit a small beep to help you find it on a crowded shelf! Such tiny, relatively dumb devices would communicate with larger, slightly less dumb ones, also scattered around the office and building. Even very familiar objects, such as the windows of a house, may gain new functionality, recording traces and trails of activity around the house. Spaces in the parking lot communicate their presence and location to the car-and-driver system via a small mirror display, and the coffee-maker in your office immediately knows when and where you have parked the car, and can prepare a hot beverage ready for your arrival.

The idea, then, is to embody and distribute the computation. Instead of focusing on making a richer and richer interface with an even more potent black box on the table, ubiquitous computing aims to make the interfaces multiple, natural, and so simple as to become rapidly invisible to the user.

The computer is thus drawn into the real world of daily objects and interactions where its activities and contributions become part of the unremarked backdrop upon which the biological brain and organism learn to depend.

This is a powerful and appealing vision. But what has it to do with the individual's status as a human-machine hybrid? Surely, I hear you saying, a smart world cannot a cyborg make. My answer: it depends just how smart the world is, and more importantly, how responsive it is, over time, to the activities and projects distinctive of an individual person. A smart world, which takes care of many of the functions that might otherwise occupy our conscious attention, is, in fact, already functioning very much like the cyborg of Clynes and Kline's original vision. The more closely the smart world becomes tailored to an individual's specific needs, habits, and preferences, the harder it will become to tell where that person stops and this tailor-made, co-evolving smart world begins. At the very limit, the smart world will function in such intimate harmony with the biological brain that drawing the line will serve no legal, moral, or social purpose. It would be as if someone tried to argue that the "real me" excludes all those nonconscious neural activities on which I so constantly depend relegating all this to a mere smart inner environment. The vision of the mind and self that remains following this exercise in cognitive amputation is thin indeed!

In what ways, then, might an electronically infested world come to exhibit the right kinds of boundary-blurring smarts? One kind of example, drawn from the realm of current commercial practice, is the use of increasingly responsive and sophisticated software agents. An example of a software agent would be a program that monitors your online reading and buying habits, and which searches out new items that fit your interests. More sophisticated software agents might monitor online auctions, bidding and selling on your behalf, or buy and sell your stocks and shares. Pattie Maes, who works on software agents at MIT media lab, describes them as

> software entities . . . that are typically long-lived, continuously running . . . and that can help you keep track of a certain task . . . so it's as if you were extending your brain or expanding your brain by having software entities out there that are almost part of you.

Reflect on the possibilities. Imagine that you begin using the web at the age of four. Dedicated software agents track and adapt to your emerging

interests and random explorations. They then help direct your attention to new ideas, web pages, and products. Over the next seventy-some years you and your software agents are locked in a complex dance of co-evolutionary change and learning, each influencing, and being influenced by, the other. You come to expect and trust the input from the agents much as you expect and trust the input from your own unconscious brain—such as that sudden idea that it would be nice to go for a drive, or to buy a Beatles CD—ideas that seem to us to well up from nowhere but which clearly shape our lives and our sense of self. In such a case and in a very real sense, the software entities look less like part of your problem-solving environment than part of you. The intelligent system that now confronts the wider world is biological-you-plus-the-software-agents. These external bundles of code are contributing as do the various nonconscious cognitive mechanisms active in your own brain. They are constantly at work, contributing to your emerging psychological profile. You finally count as "using" the software agents only in the same attenuated and ultimately paradoxical way, for example, that you count as "using" your posterior parietal cortex.

The biological design innovations that make all this possible include the provision (in us) of an unusual degree of cortical plasticity and the (related) presence of an unusually extended period of development and learning (childhood).[29] These dual innovations (intensively studied by the new research program called "neural constructivism") enable the human brain, more than that of any other creature on the planet, to factor an open-ended set of biologically external operations and resources deep into its own basic modes of operation and functioning. It is the presence of this unusual plasticity that makes humans (but not dogs, cats, or elephants) *natural-born cyborgs*: beings primed by Mother Nature to annex wave upon wave of external elements and structures as part and parcel of their own extended minds.

This gradual interweaving of biological brains with nonbiological resources recapitulates, in a larger arena, the kind of sensitive co-development found within a single brain. A human brain, as we shall later see in more detail, comprises a variety of relatively distinct, but densely intercommunicating subsystems. Posterior parietal subsystems, to take an example mentioned earlier, operate unconsciously when we reach out to grasp an object, adjusting hand orientation and finger placement appropriately.[30] The conscious agent seldom bothers herself with these details: she simply decides to reach for the object, and does so, fluently and efficiently. The conscious parts of her brain learned long ago that they could simply count on the posterior parietal structures to kick in and fine-tune the reaching as needed. In just the same way, the conscious and unconscious parts of the brain learn to factor in the operation of various nonbiological tools and resources, creating an extended problem-solving matrix whose degree of fluid integration can sometimes rival that found within the brain itself.

Let's return, finally, to the place we started: the cyborg control of aspects of the autonomic nervous system. The functions of this system (the homeostatic control of heart rate, blood pressure, respiration, etc.) were the targets of Clynes and Kline in the original 1960 proposal. The cyborg, remember, was to be a human agent with some additional, machine-controlled, layers of automatic (homeostatic) functioning, allowing her to survive in alien or inhospitable environments. Such cyborgs, in the words of Clynes and Kline, would provide "an organizational system in which such robot-like problems were taken care of automatically, leaving man free to explore, to create, to think and to feel." Clynes and Kline were adamant that such off-loading of certain control functions to artificial devices would in no way change our nature as human beings. They would simply free the conscious mind to do other work.

This original vision, pioneering though it was, was also somewhat too narrow. It restricted the imagined cyborg innovations to those serving various kinds of bodily maintenance. There might be some kind of domino effect on our mental lives, freeing up conscious neural resources for better things, but that would be all. My claim, by contrast, is that various kinds of deep human-machine symbiosis really do expand and alter the shape of the psychological processes that make us who we are. The old technologies of pen and paper have deeply impacted the shape and form of biological reason in mature, literate brains. The presence of such technologies, and their modern and more responsive counterparts, does not merely act as a convenient wrap around for a fixed biological engine of reason. Nor does it merely free up neural resources. It provides instead an array of resources to which biological brains, as they learn and grow, will *dovetail* their own activities. The moral, for now, is simply that this process of fitting, tailoring, and factoring in leads to the creation of extended computational and mental organizations: reasoning and thinking systems distributed across brain, body, and world. And it is in the operation of these extended systems that much of our distinctive human intelligence inheres.

Such a point is not new, and has been well made by a variety of theorists working in many different traditions.[31] I believe, however, that the idea of human cognition as subsisting in a hybrid, extended architecture (one which includes aspects of the brain and of the cognitive technological envelope in which our brains develop and operate) remains vastly under-appreciated. We cannot understand what is special and distinctively powerful about human thought and reason by simply paying lip service to the importance of the web of surrounding structure. Instead, we need to understand in detail how brains like ours dovetail their problem-solving activities to these additional resources, and how the larger systems thus created operate, change, and evolve. In addition, we need to understand that the very ideas of minds and persons are not limited to the biological skin-bag, and that our sense of self, place, and potential are all malleable constructs ready to expand, change, or contract at surprisingly short notice.

Consider a little more closely the basic biological case. Our brains provide both some kind of substrate for conscious thought, and a vast panoply of thought and action guiding resources that operate quite unconsciously. You do not *will* the motions of each finger and joint muscle as you reach for the glass or as you return a tennis serve. You do not *decide* to stumble upon such-and-such a good idea for the business presentation. Instead, the idea just occurs to you, courtesy once again of all those unconsciously operating processes. But it would be absurd, unhelpful, and distortive to suggest that your true nature—the real "you," the real agent—is somehow defined only by the operation of the conscious resources, resources whose role may indeed be significantly less than we typically imagine. Rather, our nature as individual intelligent agents is determined by the full set of conscious and unconscious tendencies and capacities that together support the set of projects, interests, proclivities, and activities distinctive of a particular person. Just who we are, on that account, may be as much informed by the specific sociotechnological matrix in which the biological organism exists as by those various conscious and unconscious neural events that happen to occur inside the good old biological skin-bag.

Once we take all this on board, however, it becomes obvious that even the technologically mediated incorporation of additional layers of unconscious functionality must make a difference to our sense of who and what we are; as much of a difference, at times, as do some very large and important chunks of our own biological brain. Well-fitted transparent technologies have the potential to impact what we feel capable of doing, where we feel we are located, and what kinds of problems we find ourselves capable of solving. It is, of course, *also* possible to imagine bioelectronic manipulations, which quite directly affect the contents of conscious awareness. But direct accessibility to individual conscious awareness is not essential for a human-machine merger to have a profound impact on who and what we are. Indeed, as we saw, some of the most far-reaching near-future transformations may be rooted in mergers that make barely a ripple on the thin surface of our conscious awareness.

That this should be so is really no surprise. We already saw that what we cared about, even in the case of the classic cyborgs, was some combination of seamless integration and overall transformation. But the most seamless of all integrations, and the ones with the greatest potential to transform our lives and projects, are often precisely those that operate deep beneath the level of conscious awareness. New waves of almost invisible, user-sensitive, semi-intelligent, knowledge-based electronics and software are perfectly posed to merge seamlessly with individual biological brains. In so doing they will ultimately blur the boundary between the user and her knowledge-rich, responsive, unconsciously operating electronic environments. More and more parts of our worlds will come to share the moral and psychological status of parts of our brains. We are already primed by nature to dovetail our minds to our worlds. Once the world starts dovetailing back in earnest, the last few seams must burst, and we will stand revealed: cyborgs without surgery, symbionts without sutures.

CHAPTER 2

Technologies to Bond With

Heavy Metal

Los Alamos National Laboratory, New Mexico, occupies the high ground both physically and technologically. I am here, this hot and sunny day in May 1999, to deliver a talk on the interactions between mind and technology. Getting in is not easy. There have been security scares, and my permanent resident alien card is deemed insufficient proof of identity. After a flurry of panic, my secretary somehow manages to fax them a copy of my UK passport. At last—and just in time for the talk—I am issued with the inevitable plastic photo ID. Soon I find myself deep in the radiation-proof concrete bunkers that currently serve as home to my hosts, the Complex Systems Modeling Team.

Walking around this eerily silent, windowless, underground laboratory, I am struck by the stark contrast between old technology and new. The massive concrete bunkers and reinforced floors of these old buildings were designed both to resist nuclear attack and to support heavy, in-your-face technology: giant mainframes, immense monoliths of dials, lights, and levers. Yet today's action, in the Complex Systems Laboratory at least, usually requires little more than a few potent laptops and some fiber-optic links to massive databases. The heaviest piece of real, working machinery that I encounter is a somewhat sick old printer whose wheezing vibrations occasionally disturb the tomb-like silence.