



---

Where Brain, Body, and World Collide

Author(s): Andy Clark

Source: *Daedalus*, Vol. 127, No. 2, The Brain (Spring, 1998), pp. 257-280

Published by: The MIT Press on behalf of American Academy of Arts & Sciences

Stable URL: <http://www.jstor.org/stable/20027499>

Accessed: 29/06/2009 15:40

---

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=mitpress>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We work with the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact [support@jstor.org](mailto:support@jstor.org).



The MIT Press and American Academy of Arts & Sciences are collaborating with JSTOR to digitize, preserve and extend access to *Daedalus*.

## Where Brain, Body, and World Collide

THE BRAIN FASCINATES BECAUSE IT IS the biological organ of mindfulness itself. It is the inner engine that drives intelligent behavior. Such a depiction provides a worthy antidote to the once-popular vision of the mind as somehow lying outside the natural order. But it is a vision with a price, for it has concentrated much theoretical attention on an uncomfortably restricted space—the space of the inner neural machine, divorced from the wider world that then enters the story only via the hygienic gateways of perception and action. Recent work in neuroscience, robotics, and psychology casts doubt on the effectiveness of such a shrunken perspective. Instead, it stresses the unexpected intimacy of brain, body, and world and invites us to attend to the structure and dynamics of extended adaptive systems—ones involving a much wider variety of factors and forces. While it needs to be handled with some caution, I believe there is much to be learned from this broader vision. The mind itself, if such a vision is correct, is best understood as the activity of an essentially *situated* brain: a brain at home in its proper bodily, cultural, and environmental niche.

### SOFTWARE

Humans, dogs, ferrets—these are, we would like to say, mindful things. Rocks, rivers, and volcanoes are not. And no doubt there are plenty of cases in between, such as insects or bacteria. In the

---

*Andy Clark is Professor of Philosophy and Director of the Philosophy/Neuroscience/ Psychology Program at Washington University in St. Louis.*

natural order, clear cases of mindfulness always involve creatures with brains. Hence, in part, the fascination of the brain: understanding the brain seems crucial to the project of understanding the mind. But how should such an understanding proceed?

An early sentiment—circa 1970, and no longer much in vogue—was that understanding the mind depended rather little on understanding the brain. The brain, it was agreed, was in some sense the physical medium of cognition. But everything that mattered about the brain qua mind-producing engine turned not on the physical details but on the computational and information-processing strategies that the neural stuff (merely) implemented. There was something importantly right about this view, but something desperately wrong as well.

What was right was the observation that understanding the physical workings of the brain was not going to be enough; we would need also to understand how the system was organized at a higher level in order to grasp the roots of mindfulness in the firing of neurons. This point is forcefully made by the cognitive scientist Brian Cantwell Smith, who draws a parallel with the project of understanding ordinary computer systems. Consider, for example, a standard personal computer running a tax-calculation program. We could quite easily answer all the “physiological” questions (using source code and wiring diagrams) while still lacking any real understanding of what the program does or even how it works.<sup>1</sup> To really understand how mental activity yields mental states, many theorists believe, we must likewise understand something of the computational/information-processing organization of the brain. Physiological studies may contribute to this understanding. But even a full physiological story would not, in and of itself, reveal how brains work qua mind-producing engines.<sup>2</sup>

The danger, of course, was that this observation could be used as an excuse to downplay or marginalize the importance of looking at the biological brain *at all*. And so it was that, in the early days of cognitive science, it was common to hear real neuroscience dismissed as having precious little to offer to the general project of understanding human intelligence. Such a dismissal, however, could not long be sustained. For although it

is (probably) true to say that a computational understanding is in principle independent of the details of any specific implementation in hardware (or wetware),<sup>3</sup> the project of *discovering* the relevant computational description (especially for biological systems) is surely not.

One key factor is evolution. Biological brains are the product of biological evolution and as such often fail to function in the ways we (as human designers) might expect.<sup>4</sup> This is because evolution is both constrained and liberated in ways we are not. It is constrained to build its solutions incrementally via a series of simpler but successful ancestral forms. The human lung, to give one example, is built via a process of tinkering with the swim bladder of the fish.<sup>5</sup> The engineer might design a better lung from scratch. The tinkerer, by contrast, must take an existing device and subtly adapt it to a new role. From the engineer's ahistorical perspective, the tinkerer's solution may look bizarre. Likewise, the processing strategies used by biological brains may surprise the computer scientist. Such strategies have themselves been evolved via a process of incremental, piecemeal tinkering with older solutions.

More positively, biological evolution is *liberated* by being able to discover efficient but "messy" or unobvious solutions, ones that may, for example, exploit environmental interactions and feedback loops so complex that they would quickly baffle a human engineer. Natural solutions (as we shall later see) can exploit just about any mixture of neural, bodily, and environmental resources along with their complex, looping, and often nonlinear interactions. Biological evolution is thus able to explore a very different solution space (wider in some dimensions, narrower in others) than that which beckons to conscious human reason.

There are, of course, ways around this apparent mismatch. Computationalists lately exploit so-called "genetic algorithms" that roughly mimic the natural process of evolutionary search and allow the discovery of efficient but loopy and interactive adaptive strategies.<sup>6</sup> Moreover, hard neuroscientific data and evidence is increasingly available as a means of helping expand our imaginative horizons so as to appreciate better the way real biological systems solve complex problems. The one-time obses-

sion with the “software level” is thus relaxing, in favor of an approach that grounds computational modeling in a more serious appreciation of the biological facts. In the next section I highlight some challenging aspects of such recent research— aspects that will lead us directly to a confrontation with the situated brain.

#### WETWARE, AND SOME ROBOTS

Recent work in cognitive neuroscience underlines the distance separating biological and “engineered” solutions to problems, and it displays an increasing awareness of the important interpenetration—in biological systems—of perception, thought, and action. Some brief examples should help fix the flavor.

As a gentle entry point, consider some recent work on the neural control of monkey finger motions. Traditional wisdom depicted the monkey’s fingers as individually controlled by neighboring groups of spatially clustered neurons. According to this story, the neurons (in motor area 1, or M1) were organized as a “somatotopic map” in which a dedicated neural subregion governed each individual digit, with the subregions arranged in lateromedial sequence just like the fingers on each hand. This is a tidy, easily conceptualized solution to the problem of finger control. But it is the engineer’s solution, not (it now seems) that of nature.

Marc Schieber and Lyndon Hibbard have shown that individual digit movements are accompanied by activity spread pretty well throughout the M1 hand area, and that precise, single-digit movements actually require *more* activity than some multidigit whole hand actions (such as grasping an object).<sup>7</sup> Such results are inconsistent with the hypothesis of digit-specific local neuronal groups. From a more evolutionary perspective, however, the rationale and design is less opaque. Some time earlier Schieber had conjectured that the basic case, from an evolutionary perspective, is the case of whole-hand grasping motions (used to grab branches, swing, acquire fruits, and so forth), and that the fundamental neural adaptations are thus geared to the use of simple commands that exploit inbuilt synergies of muscle and tendon so as to yield such coordinated motions.<sup>8</sup> The “complex”

coordinated case is thus evolutionarily basic and neurally atomic. The “simple” task of controlling, for example, an individual digit represents the harder problem and requires more neural activity—viz., the use of some motor cortex neurons to *inhibit* the synergetic activity of the other digits. Precise single-digit movements require the agent to tinker with whole-hand commands, modifying the basic synergetic dynamics (of mechanically linked tendons, etc.) adapted to the more common task.

Consider next a case of perceptual adaptation. The human perceptual system can, we know—given time and training—adapt in quite powerful ways to distorted or position-shifted inputs. For example, subjects can learn how to coordinate vision and action while wearing lenses that invert the entire visual scene so that the world initially appears upside down. After wearing such lenses for a few days, the world is seen to flip over—various aspects of the world now appear to the subject to be in the normal upright position. Remove the lenses and the scene is again inverted until readaptation occurs.<sup>9</sup> William Thach and his colleagues at Washington University used a variant of such experiments to demonstrate the motor-specificity of some perceptual adaptations. Wearing lenses that shifted the scene *sideways* a little, subjects were asked to throw darts at a board. In this case, repeated practice led to successful adaptation, but of a motor loop-specific kind.<sup>10</sup> The compensation did not “carry over” to tasks involving the use of the nondominant hand to throw or to an underarm variant of the throw. Instead, adaptation looked to be restricted to a quite-specific combination of gaze angle and throwing angle: the one used in overarm, dominant-hand throwing.

Something of the neural mechanisms of such adaptation is now understood. It is known, for example, that the adaptation never occurs in patients with generalized cerebellar cortical atrophy, and that inferior olive hypertrophy leads to impaired adaptation. On the basis of this and other evidence, Thach and his coworkers speculated that a learning system implicating the inferior olive and the cerebellum (linked via climbing fibers) is active both in prism adaptation and in the general learning of patterned responses to frequently encountered stimuli. The more general lesson, however, concerns the nature of the perception-

action system itself. It increasingly appears that the simple image of a general-purpose perceptual system delivering input to a distinct and fully independent action system is biologically distortive. Instead, perceptual and action systems work together, in the context of specific tasks, so as to promote adaptive success. Perception and action, on this view, form a deeply interanimated unity.

Further evidence for such a view comes from a variety of sources. Consider, for example, the fact that the primate visual system relies on processing strategies that are not strictly hierarchic but instead depend on a variety of top-to-bottom and side-to-side channels of influence. These complex inner pathways allow a combination of multiple types of information (high-level intentions, low-level perception, and motor activity) to influence all stages of visual processing. The macaque monkey (to take one well-studied example) possesses about thirty-two visual brain areas and over three hundred connecting pathways. The connecting pathways go both upwards and downwards (for example, from V1 to V2 and back again) and side to side (between subareas in V1).<sup>11</sup> Individual cells at “higher” levels of processing, such as visual area 4 (V4), do, it is true, seem to specialize in the recognition of specific geometric forms. But they will each also respond, to some small degree, to many other stimuli. These small responses, spread unevenly across a whole population of cells, can carry significant information. The individual cells thus function not as narrowly tuned single-feature detectors but as widely tuned filters reacting to a whole range of stimulus dimensions.<sup>12</sup> Moreover, the responses of such cells now look to be modifiable both by attention and by details of local task-specific context.<sup>13</sup>

More generally, back-projecting (corticocortical) connections tend, in the monkey, to outnumber forward-projecting ones; in other words, there are more pathways leading from deep inside the brain outwards towards the sensory peripheries than vice versa.<sup>14</sup> Visual processing may thus involve a variety of criss-crossing influences that could only roughly, if at all, be described as a neat progression through a lower-to-higher hierarchy.

Such complex connectivity opens up a wealth of organizational possibilities in which multiple sources of information

combine to support visually guided action. Examples of such combinations are provided by Patricia Churchland, V. S. Ramachandran, and Terrence Sejnowski, who offer a neuro-physiologically grounded account of what they term “interactive vision.”<sup>15</sup> The interactive vision paradigm is there contrasted with approaches that assume a simple division of labor in which perceptual processing yields a rich, detailed inner representation of the visual scene—one that is later given as input to the reasoning and planning centers, which in turn calculate a course of action and send commands to the motor effectors. This simple image (of what roboticists call a “sense-think-act” cycle) is, it now seems, not true to the natural facts. In particular, daily agent-environment interactions often do not require the construction and use of detailed inner models of the full visual scene. Moreover, low-level perception may “call” motor routines that yield *better perceptual input* and hence improve information pickup. Beyond this, real-world actions may sometimes play an important role in the computational process itself. Finally, the internal representation of worldly events and structures may be less like a passive data structure or description and more like a direct recipe for action.

Evidence for the first of these propositions comes from a series of experiments in which subjects watch pictorial images on a computer screen. Then, as they continue to saccade around the scene (focusing first on one area, then another), small changes are made to the currently unattended parts of the display. The changes are made during the visual saccades. It is an amazing fact that, for the most part, quite large changes go unnoticed: changes such as the replacement of a tree by a shrub, or the addition of a car, the deletion of a hat, and so on.<sup>16</sup> Why do such gross alterations remain undetected? A compelling hypothesis is that the visual system is not even attempting to build a rich, detailed model of the current scene but is instead geared to using frequent saccades to retrieve information *as and when it is needed* for some specific problem-solving purpose. This fits nicely with Yarbus’s classic finding that the pattern of such saccades varies (even with identical scenes) according to the type of task the subject has been set (for example, to give the ages of the people in a picture, to guess the activity they have been engaged

in, and so on).<sup>17</sup> One explanation of our subjective impression of forming a rich inner representation of the whole scene is that we are prone to this illusion because we are able to perform fast saccades, retrieving information as and when required. (An analogy: a modern store may present the illusion of having a massive amount of goods stocked on the premises, because it always has what you want when you want it. But modern computer ordering systems can automatically count off sales and requisition new items so that the necessary goods are available just when needed and barely a moment before. This fine-tuned ordering system offers a massive saving of on-site storage while tailoring supply directly to customer demand.)<sup>18</sup>

Contemporary research in robotics avails itself of these same economies. Rodney Brooks, one of the pioneers of “new robotics,” coined the slogan “the world is its own best model” to capture just this flavor.<sup>19</sup> A robot known as Herbert, to take just one example, was designed to collect soft-drink cans left around a crowded laboratory. But instead of requiring powerful sensing capacities and detailed advance planning, Herbert got by very successfully using a collection of coarse sensors and simple, relatively independent, behavioral routines. Basic obstacle avoidance was controlled by a ring of ultrasonic sensors that brought the robot to a halt if an object was in front of it. General locomotion (randomly directed) was interrupted if Herbert’s simple visual system detected a rough tablelike outline. At this point a new routine kicks in and the table surface is swept using a laser. If the outline of a can is detected, the whole robot rotates until the can is centered in its field of vision. This simple physical action simplifies the pickup procedure by creating a standard action-frame in which a robot arm, equipped with simple touch sensors, gently skims the table surface dead ahead. Once a can is encountered, it is grasped and collected, and the robot moves on. Notice, then, that Herbert succeeds without using any conventional planning techniques and without creating and updating any detailed inner model of the environment. Herbert’s “world” is composed of undifferentiated obstacles and rough tablelike and canlike outlines. Within this world the robot also exploits its own bodily actions (rotating the “torso” to center the can in its field of view) so as to simplify greatly the computational

problem involved in eventually reaching for the can. Herbert is thus a simple example of the second proposition above—a system that succeeds using minimal representational resources, and one in which gross motor activity helps streamline a perceptual routine.<sup>20</sup>

The interactive vision framework envisages a more elaborate natural version of this same broad strategy, viz., the use of a kind of perceptuomotor loop whose role is to make the most of incoming perceptual information by combining multiple sources of information. The idea here is that perception is not a passive phenomenon through which motor activity is initiated only at the endpoint of a complex process whereby the animal creates a detailed representation of the perceived scene. Instead, perception and action engage in a kind of incremental game of tag, in which motor assembly begins long before sensory signals reach the top level. Thus, early perceptual processing may yield a kind of protoanalysis of the scene, enabling the creature to select actions (such as head and eye movements) whose role is to provide a slightly upgraded sensory signal. That signal may, in turn, yield a new protoanalysis indicating further visuomotor action, and so on. Even whole-body motions may be deployed as part of this process of improving perceptual pickup. Foveating an object can, for example, involve motion of the eyes, head, neck, and torso. Churchland and her associates put it well: “Watching Michael Jordan play basketball or a group of ravens steal a caribou corpse from a wolf tends to underscore the integrated, whole-body character of visuomotor coordination.”<sup>21</sup> This integrated character is consistent with neurophysiological and neuroanatomical data that show the influence of motor signals in visual processing. There are—to take just two small examples—neurons sensitive to eye position in V1, V3, and LGN (lateral geniculate nucleus), and cells in V1 and V2 that seem to know in advance about planned visual saccades (i.e., they show enhanced sensitivity to the target).<sup>22</sup>

Moving on to the third proposition—that real-world actions may sometimes play an important role in the computational process itself—consider the task of distinguishing a figure from the ground (the rabbit from the field, or whatever). It turns out that this problem is greatly simplified using information ob-

tained from head movement during eye fixation. Likewise, depth perception is greatly simplified using cues obtained by the observer's own self-directed motion. As the observer moves, close objects will show more relative displacement than farther ones. That is probably why, as Churchland, Ramachandran, and Sejnowski observe, head-bobbing behavior is frequently seen in animals: "A visual system that integrates across several glimpses to estimate depth has computational savings over one that tries to calculate depth from a single snapshot."<sup>23</sup>

And so, with the last proposition—that the neural representation of worldly events may be less like a passive data structure and more like a recipe for action—the driving force, once again, is computational economy. If the goal of perception and reason is to guide action (and it surely is, evolutionarily speaking), it will often be simpler to represent the world in ways rather closely geared to the kinds of actions we want to perform. To take a simple example, an animal that uses its visual inputs to guide a specific kind of reaching behavior (so as to acquire and ingest food) need not form an object-centered representation of the surrounding space. Instead, a systematic metrical transformation (achieved by a point-to-point mapping between two topographic maps) may transform the visual inputs directly into a recipe for reaching out and grabbing the food. In such a setup, the animal does not need to do any computational work on an action-neutral inner model so as to plan a reaching trajectory. The perceptual processing is instead tweaked, at an early stage, in a way dictated by the particular use to which the visual input is dedicated. This strategy is described in detail in Paul Churchland's account of the "connectionist crab," in which research in artificial neural networks is applied to the problem of creating efficient point-to-point linkages between deformed topographic maps.<sup>24</sup>

In a related vein, Maja Mataric of the Massachusetts Institute of Technology's Artificial Intelligence Laboratory has developed a neurobiologically inspired model of how rats navigate their environments. This model exploits the kind of layered architecture also used in the robot Herbert.<sup>25</sup> Of most immediate interest, however, is the way Mataric's robot learns about its surroundings. As it moves around a simple maze, it detects land-

marks that are registered as a combination of sensory input and current motion. A narrow corridor thus registers as a combination of forward motion and short lateral distance readings from sonar sensors. Later, if the robot is required to find its way back to a remembered location, it retrieves an interlinked body of such combined sensory and motor readings.<sup>26</sup> The stored “map” of the environment is thus immediately fit to act as a recipe for action, since the motor signals are part of the stored knowledge. The relation between two locations is directly encoded as the set of motor signals that moved the robot from one to the other. The inner map is thus *itself* the recipe for the necessary motor actions. By contrast, a more classical approach would first generate a more objective map, which would then need to be *reasoned over* in order to plan the route.

The Mataric robot and the connectionist crab exemplify the attractions of what I call “action-oriented representations,” representations that describe the world by depicting it in terms of possible actions.<sup>27</sup> This image fits in nicely with several of the results reported earlier, including the work on monkey finger control and the motor-loop specificity of “perceptual” adaptation. The products of perceptual activity, it seems, are not always action-neutral descriptions of external reality. They may instead constitute direct recipes for acting and intervening. We thus glimpse something of the shape of what some have described as a framework that is “motocentric” rather than “visuocentric.”<sup>28</sup>

As a last nod in this same direction, consider the fascinating case of so-called mirror neurons.<sup>29</sup> These are neurons found in the monkey ventral premotor cortex that are action-oriented, context-dependent, and implicated in both self-initiated activity and passive perception. They are active both when the monkey observes a specific action (such as someone grasping a food item) and when the monkey performs the same action, where sameness implies not mere grasping but the grasping of a food item.<sup>30</sup> The implication, according to the psychologist and neuroscientist Marc Jeannerod, is that “the action. . .to be initiated is stored in terms of an action code, not a perceptual one.”<sup>31</sup>

Putting all this together suggests a much more integrated model of perception, cognition, and action. Perception is itself

tangled up with possibilities for action and is continuously influenced by cognitive, contextual, and motor factors. It need not yield a rich, detailed, and action-neutral inner model awaiting the services of “central cognition” so as to deduce appropriate actions. In fact, these old distinctions (between perception, cognition, and action) may sometimes obscure, rather than illuminate, the true flow of effect. In a certain sense, the brain is revealed not as (primarily) the engine of reason or quiet deliberation, but as the organ of *environmentally situated control*. Action, not truth and deductive inference, is the key organizing concept. This perspective, however, begs to be taken a step further. For if brains are best understood as controllers of environmentally situated activity, then might it not be fruitful to locate the neural contribution as just one (important) element in a complex causal web, spanning brains, bodies, and world? This potential gestalt shift is the topic of the next section.

#### WIDEWARE

Let us coin a term, “wideware,” to refer to states, structures, or processes that satisfy two conditions. First, the item in question must be in some intuitive sense environmental; it must not, at any rate, be realized within the biological brain or the central nervous system. Bodily aspects and motions, as well as truly external items such as notebooks and calculators, thus fit the bill. Second, the item (state, structure, process) must play a functional role as part of an *extended cognitive process*: a process geared to the promotion of adaptive success via the gathering and use of knowledge and information, and one that loops out in some nontrivial way, so as to include and exploit aspects of the local bodily and environmental setting.

Of course, even what is intuitively a fully internalized cognitive process will usually involve contact with the external world. That much is demanded by the very ideas of knowledge acquisition and information-gathering. The notion of wideware aims to pick out instead a somewhat more restricted range of cases—ones in which the *kind* of work that cognitive science has typically assigned to the inner workings of the brain is at least partly carried out by processes of storage, search, and transformation,

realized using bodily actions and/or a variety of external media. Understanding the human mind, I shall argue, will require us to attend much more closely to the role of such bodily actions and external media than was once anticipated.

To better fix the notion of *wideware* itself, consider first a very simple case involving “bodily backdrop.” Computer-controlled machines are sometimes used to fit small parts into one another. The error tolerance here is very low; sometimes the robot arm will fail to make a match. A computationally expensive solution exists in which the control system includes multiple feedback loops that signal such failures and prompt the machine to try again in a minutely different orientation. It turns out, however, that a much cheaper, more robust, and more efficient procedure is simply to mount the assembler arms on rubber joints that “give” along two spatial axes. This bodily backdrop allows the control device to dispense with the complex feedback loops. Thanks to the rubber mountings, the parts “jiggle and slide into place just as if millions of tiny feedback adjustments to a rigid system were being continuously computed.”<sup>32</sup>

Mere bodily backdrop, however, does not really count as an instance of an extended cognitive process. The computational and information-based operations are reduced, courtesy of the brute physical properties of the body. But they are not themselves extended *into* the world. Genuine cognitive and computational extension requires, by contrast, that the external or bodily operations are themselves usefully seen as performing cognitive or information-processing operations. A simple phototropic (light-following) robot—to take another negative example—does not constitute an extended cognitive or computational system. Although the presence of some external structure (light sources) is here vital to the robot’s behavioral routines, those external structures are not doing the kind of work that cognitive science and psychology has typically assigned to inner neural activity. They are not acting so as to manipulate, store, or modify the knowledge and information that the organism uses to reach its goals. Sometimes, however, external structures and bodily operations do seem to be proper parts of the cognitive and computational processes themselves. Thus consider the use, in recent interactive vision research, of so-called deictic pointers. A classical (nondeictic)

pointer is an inner state that can figure in computational routines and that can also “point” to further data structures. Such pointing allows for both the retrieval of additional information and the binding of the content of one memory location to another. One use of classical pointers is thus to bind information about spatial location (“top left corner of the visual field”) to information about current features (“bright yellow mug”), yielding complex contents (in this case, “bright yellow mug in the top left corner of the visual field”).<sup>33</sup> Since neural processing involves more or less distinct channels for object properties, object location, and object motion (for example), binding is an important element in neural computation. But binding, it now seems, can sometimes be achieved by the use of actual bodily orientations instead of linking inner-data structures.

The story is complex, but the basic idea is straightforward. It is to set up a system so that bodily orientations (such as saccadic eye motions leading to object fixation) directly yield the kinds of benefits associated with classical binding. Dana Ballard and colleagues have shown how to use visual scene fixation to directly associate the features represented (the properties of the perceived object) with an external spatial location.<sup>34</sup> Another example is the use of so-called “do-it-where-I’m-looking” processing routines in which a bodily motion (say, grasping) is automatically directed to whatever location happens to be currently fixated in the visual field. In all these cases, the authors comment: “The external world is analogous to computer memory. When fixating a location the neurons that are linked to the fovea refer to information computed from that location. Changing gaze is analogous to changing the memory reference in a silicon computer.”<sup>35</sup>

Deictic binding provides a clear example of a case in which bodily motion does the *kind* of work that cognitive science has typically assigned only to inner neural activity. The deictic strategies use a combination of inner computation and gross bodily action to support a type of functionality once studied as an exclusive property of the inner neural system. Taking this perspective a step further, we next consider the use of external media as both additional memory and as potent symbol-manipulating arenas.

Portions of the external world, it is fairly obvious, often function as a kind of extraneural memory store. We may deliberately leave a film on our desk to remind us to take it for developing. Or we may write a note (“develop film”) on paper and leave that on our desk instead. As users of words and texts, we command an especially potent means of off-loading data and ideas from the biological brain onto a variety of external media. This trick, I think, is not to be underestimated, for it affects not just the quantity of data at our command but also the kinds of operation we can bring to bear on it. Words, texts, symbols, and diagrams often figure *deeply* in the problem-solving routines developed by biological brains nurtured in language-rich environmental settings. Human brains, trained in a sea of words and text, will tend to develop computational strategies that directly factor in the reliable presence of a variety of external props and aids. The inner operations will then complement, but not replicate, the special manipulative potentials provided by the external media.

Consider, for example, the process of writing an academic paper.<sup>36</sup> You work long and hard, and at the end of the day you are happy. Being a good physicalist, you assume that all the credit for the final intellectual product belongs to your brain, the seat of human reason. But you are too generous by far. What really happened was (perhaps) more like this: The brain supported some rereading of old texts, materials, and notes. While rereading these, it responded by generating a few fragmentary ideas and criticisms. These ideas and criticisms were then stored as more marks on paper, in margins, on computer disks, or some other recording means. The brain then played a role in reorganizing this data on clean sheets, adding new on-line reactions and ideas. The cycle of reading, responding, and external reorganization is repeated, again and again. Finally, there is a product—a story, argument, or theory. But this intellectual product owes a lot to those repeated loops into the environment. Credit belongs to the agent-in-the-world. The biological brain is just a part (albeit a crucial and special part) of a spatially and temporally extended process, involving lots of extraneural operations, whose joint action creates the intellectual product. There is thus a real sense (or so I would argue) in which the notion of the

“problem-solving engine” is really the notion of the *whole caboodle*: the brain and body operating within an environmental setting.

Consider, by way of analogy, the idea of a swimming machine; in particular, consider the bluefish tuna.<sup>37</sup> The tuna is paradoxically talented. Physical examination suggests it should not be able to achieve the aquatic feats of which it is demonstrably capable. It is physically too weak (by about a factor of 7) to swim as fast as it does, to turn as compactly as it does, or to move off with the acceleration it does. The explanation (according to the fluid dynamicists Michael and George Triantafyllou) is that these fish actively create and exploit additional sources of propulsion and control in their watery environments. For example, the tuna use naturally occurring eddies and vortices to gain speed, and they flap their tails so as actively to create additional vortices and pressure gradients that they then exploit for quick takeoffs and similar feats. The real swimming machine, I suggest, is thus the fish *in its proper context*: the fish plus the surrounding structures and vortices than it actively creates and then maximally exploits. The *cognitive machine*, in the human case, looks similarly extended.<sup>38</sup> We actively create and exploit multiple linguistic media, yielding a variety of contentful structures and manipulative opportunities whose reliable presence is then factored deep into our problem-solving strategies.

#### IMPLICATIONS

Software, wetware, and wideware, if our story is to be believed, form a deeply interanimated triad. The computational activities of the brain will be heavily sculpted by its biological “implementation.” And there will be dense complementarity and cooperation between neural, bodily, and environmental forces and factors. What the brain does will thus be precisely fitted to the range of complementary operations and opportunities provided by bodily structure, motion, and the local environment. In the special case of human agency, this includes the humanly generated “whirlpools and vortices” of external, symbol-laden media: the explosion of wideware made available by the ubiquitous

devices of language, speech, and text. The picture that emerges is undeniably complex. But what does it really mean both for our understanding of ourselves and for the practice of scientific inquiry about the mind?

Certain implications for our vision of ourselves are clear. We must abandon the image of ourselves as essentially disembodied reasoning engines. And we must do so not simply by insisting that the mental is fully determined by the physical, but by accepting that we are beings whose neural profiles are profoundly geared so as to press maximal benefit from the opportunities afforded by bodily structure, action, and environmental surroundings. Biological brains are, at root, controllers of embodied action. Our cognitive profile is *essentially* the profile of an embodied and situated organism. Just how far we should then press this notion of cognitive extension, however, remains unclear. Should we just think of ourselves as cognitive agents who co-opt and exploit surrounding structures (say, pen and paper) so as to expand out problem-solving capacities? Or is there a real sense in which the cognitive agent (as opposed to the bare biological organism) is thus revealed as an extended entity incorporating brain, body, and some aspects of the local environment? Normal usage would seem to favor the former. But the more radical interpretation is not as implausible as it may initially appear.

Consider an example. Certain Alzheimer's sufferers maintain an unexpectedly high level of functioning within the normal community. These individuals should not—given their performance on a variety of standard tests—be capable of living as independently as they do. Their unusual success is explained only when they are observed in their normal home environments, in which an array of external props and aids turns out to serve important cognitive functions.<sup>39</sup> Such props and aids may include the extensive use of labels (on rooms, objects, and so forth), the use of a “memory book” containing annotated photos of friends and relatives, the use of a diary for tasks and events, and simple tactics such as leaving all important objects (say, one's checkbook) in open view so as to aid retrieval when needed.<sup>40</sup> The upshot, clearly, is an increased reliance on various forms of wideware (or “cognitive scaffolding”) as a means of

counterbalancing a neurally based deficit.<sup>41</sup> But the pathological nature of the case is, in a sense, incidental. Imagine a whole community whose linguistic and cultural practices evolved so as to counterbalance a normal cognitive profile (within that community) identical to that of these individuals. The external props could there play the same functional role (of complementing a certain neural profile) but without any overtone of pathology-driven compensation. Finally, reflect that our own community is just like that. Our typical neural profile is different, to be sure, but *relative to that profile* the battery of external props and aids—laptops, filofaxes, texts, compasses, maps, slide rules—plays just the same role. They offset cognitive limitations built into the basic biological system.

Now ask yourself what it would mean, in the case of the Alzheimer's sufferer, to maliciously damage that web of external cognitive support. Such a crime has, as Daniel Dennett once remarked, much of the flavor of a harm to the *person*, rather than simple harm to property. The same may well be true in the normal case: deliberate theft of the poet's laptop is a very special kind of crime. Certain aspects of the external world, in short, may be so integral to our cognitive routines as to count as *part of the cognitive machinery* itself (just as the whorls and vortices are, in a sense, part of the swimming machinery of the tuna). It is thus something of a question whether we should see the cognizer as the bare biological organism (that exploits all those external props and structures), or as the organism-plus-wideware. To adopt the latter perspective is to opt for a kind of "extended phenotype" view of the mind, in which the relation between the biological organism and the wideware is as important and intimate as that of the spider and the web.<sup>42</sup> \*

The implications for the specific study of the mind are, fortunately, rather less ambiguous. The vision of the brain as a controller of embodied and situated action suggests the need to develop new tools and techniques capable of investigating the brain (literally) in action—playing its part in problem-solving routines in (as far as possible) a normal ecological context. Very significant progress has already been made, of course. Noninvasive scanning techniques, such as positron emission tomography (PET) scanning and magnetic resonance imaging (MRI), represent a

giant leap beyond the use of single-cell recordings from anaesthetized animals. But we should not underestimate the distance that remains to be covered. For example, experimenters recently carried out some neural recordings from a locust in free flight.<sup>43</sup> This kind of ecologically realistic study is clearly mandated by the kinds of consideration we have put forward. Yet such investigations remain problematic due to sheer technological limitations. In the locust case, the researchers relied upon tiny radio transmitters implanted in the insect. But the information pickup remained restricted to a scanty two channels at a time.

Moreover, it is not just the information-gathering techniques that need work; the analytic tools we bring to bear on the information gathered need improving as well. If we take seriously the notion that brain, body, and world are often united in an extended problem-solving web, we may need to develop analytic and explanatory strategies that better reflect and accommodate this dense interanimation. To this end, there are (so far) two main proposals on the table. One proposal is to use the tools of dynamical systems theory. This is a well-established mathematical framework for studying the temporal evolution of states within complex systems.<sup>44</sup> Such a framework appeals, in part, because it allows us to use a single mathematical formalism to describe both internal and external organizations and (hence) allows us to treat complex looping interactions (ones that criss-cross brain, body, and world) in a deeply unified manner. The other proposal is to take the kinds of analysis we traditionally applied only to the inner states and extend it outwards.<sup>45</sup> This means sticking with talk of representations and computations but applying these ideas to extended organizations encompassing, for example, multiple individuals, maps, texts, and social institutions. (My own view, which I will not defend here, is that we need to combine the two approaches by defining new, dynamical, process-based ways of understanding key terms like “representation” and “computation.”<sup>46</sup>)

Finally, the pervasive notion of a neural code or codes is now in need of major overhaul. If the notion is to apply to real biological systems, it must be relieved of a good deal of excess baggage. Natural neural codes, as we saw above, will often be closely geared to the particular details of body and world. The

coding for single-finger motion is unexpectedly baroque, courtesy of the need to combat basic synergies created by the system of mechanically linked tendons. Other codes (for example, for whole-hand grasping motions) prove unexpectedly simple and direct. Moreover, there is evidence of the widespread use of task-specific, motor-oriented, and context-dependent encodings. Mirror neurons code for context-specific actions (say, grasping food) and function both in passive perceptions and active grasping. And interactive vision routines press large benefits from minimal and often task-specific forms of internal encoding. In all these cases, we discern a much more austere and action-oriented vision of neural encoding—a vision stripped of the excess baggage of a single, rich, languagelike inner code.<sup>47</sup>

The simple, almost commonsensical notion of the brain as a system that evolved in order to guide the actions of embodied agents in rich real-world settings thus yields substantial and sometimes challenging fruit. Gone is the vision of the environment as simply a source of problem-specifying inputs and an arena for action. Instead, both basic and imposed aspects of environmental order and complexity now emerge as fundamental components of natural problem-solving behavior. Gone too is the vision of the human brain as an organ of pure reason. In its place, we encounter the brain as a locus of action-oriented and activity-exploiting problem-solving techniques, and as a potent generator and exploiter of cognition-enhancing wideware. Taking this vision seriously, and turning it into a concrete and multidisciplinary research program, presents an exciting new challenge for the sciences of the mind.<sup>48</sup>

## ENDNOTES

<sup>1</sup>See Brian Cantwell Smith, *On the Origin of Objects* (Cambridge, Mass.: MIT Press, 1996), 148. Note that Smith's worry, at root, concerns the gap between physiological and semantic or intentional questions.

<sup>2</sup>Thus we read, for example, that computational approaches make possible "a science of structure and function divorced from material substance [that] . . . can answer questions traditionally posed by psychologists." Zenon W. Pylyshyn, ed., *The Robot's Dilemma: The Frame Problem in Artificial Intelligence* (Norwood, U.K.: Ablex, 1987), 68.

- <sup>3</sup>For example, see David Marr's distinction between the levels of computation, algorithm, and implementation. David Marr, *Vision* (San Francisco, Calif.: W. H. Freeman, 1982). For a powerful critique of the tendency of cognitive science to discount the importance of neuroscience, see George N. Reeke, Jr., and Gerald M. Edelman, "Real Brains and Artificial Intelligence," *Dædalus* 117 (1) (Winter 1988): 143–173.
- <sup>4</sup>See, for example, Herbert Alexander Simon, *Models of Bounded Rationality*, 2 vols. (Cambridge, Mass.: MIT Press, 1982); Richard Dawkins, *The Extended Phenotype: The Gene as the Unit of Selection* (Oxford: Oxford University Press, 1982); and Andy Clark, *Being There: Putting Brain, Body, and World Together Again* (Cambridge, Mass.: MIT Press, 1997), chap. 5.
- <sup>5</sup>François Jacob, "Evolution and Tinkering," *Science* 196 (1977): 1161–1166.
- <sup>6</sup>For a review, see Clark, *Being There*, chap. 5.
- <sup>7</sup>Marc Schieber and Lyndon Hibbard, "How Somatotopic is the Motor Cortex Hand Area?" *Science* 261 (1993): 489–492.
- <sup>8</sup>Marc Schieber, "How might the Motor Cortex Individuate Movements?" *Trends in Neuroscience* 13 (11) (1990): 440–444. The notion of synergy aims to capture the idea of links that constrain the collective unfolding of a system comprising many parts. For example, the front wheels of a car exhibit a built-in synergy that allows a single driver "command" (at the steering wheel) to affect them both at once. Synergetic links may also be learned, as when we acquire an automated skill, and may be neurally as well as brute-physiologically grounded. See J. A. Scott Kelso, *Dynamic Patterns* (Cambridge, Mass.: MIT Press, 1995), 38, 52.
- <sup>9</sup>For a survey of such experiments, see Robert B. Welch, *Perceptual Modification: Adapting to Altered Sensory Environments* (New York: Academic Press, 1978).
- <sup>10</sup>In this case, *without* any perceived shift in the visual scene. See W. Thach, H. Goodkin, and J. Keating, "The Cerebellum and the Adaptive Coordination of Movement," *Annual Review of Neuroscience* 15 (1992): 403–442.
- <sup>11</sup>See, for example, Daniel Felleman and David C. Van Essen, "Distributed Hierarchical Processing in the Primate Visual Cortex," *Cerebral Cortex* 1 (1991): 1–47.
- <sup>12</sup>See David C. Van Essen and Jack Gallant, "Neural Mechanisms of Form and Motion Processing in the Primate Visual System," *Neuron* 13 (1994): 1–10.
- <sup>13</sup>James Knierim and David C. Van Essen, "Visual Cortex: Cartography, Connectivity, and Concurrent Processing," *Current Opinion in Neurobiology* 2 (1992): 150–155. In a somewhat related vein, Caminiti and his coworkers showed that the directional preference of individual cells encoding reaching movement commands varies according to the initial arm position. R. Caminiti et al., "Shift of Preferred Directions of Premotor Cortical Cells with Arm Movements Performed Across the Workspace," *Experimental Brain Research* 82 (1990): 228–232. See also the discussion in Marc Jeannerod, *The Cognitive Neuroscience of Action* (Oxford: Blackwell, 1997), chap. 2.

- <sup>14</sup>Though much of the connectivity is reciprocal. See David Van Essen and Charles Anderson, "Information Processing Strategies and Pathways in the Primate Retina and Visual Cortex," in Stephen Zornetzer, Joel L. Davis, and Clifford Lau, eds., *An Introduction to Neural and Electronic Networks* (New York: Academic Press, 1990). See also Patricia Churchland, V. S. Ramachandran, and Terence Sejnowski, "A Critique of Pure Vision," in Christof Koch and Joel Davis, eds., *Large-Scale Neuronal Theories of the Brain* (Cambridge, Mass.: MIT Press, 1994), 40.
- <sup>15</sup>Churchland, Ramachandran, and Sejnowski, "A Critique of Pure Vision."
- <sup>16</sup>The exception is if subjects are told in advance to watch out for changes to a certain feature. G. W. McConkie, "Where Vision and Cognition Meet," paper presented at the H.F.S.P. Workshop on Object and Sense Perception, Leuven, Belgium, 1990. See also Churchland, Ramachandran, and Sejnowski, "A Critique of Pure Vision."
- <sup>17</sup>Alfred L. Yarbus, *Eye Movements and Vision* (New York: Plenum Press, 1967).
- <sup>18</sup>See Churchland, Ramachandran, and Sejnowski, "A Critique of Pure Vision," and Dana H. Ballard, "Animate Vision," *Artificial Intelligence* 48 (1991): 57–86. My thanks to David Clark for pointing out the store analogy.
- <sup>19</sup>See, for example, Rodney Brooks, "Intelligence Without Representation," *Artificial Intelligence* 47 (1991): 139–159.
- <sup>20</sup>Jonathan Connell, "A Colony Architecture for an Artificial Creature," *MIT AI Lab Tech Report 2* (5) (1991): 1.
- <sup>21</sup>Churchland, Ramachandran, and Sejnowski, "A Critique of Pure Vision," 44.
- <sup>22</sup>R. Wurtz and C. Mohler, "Enhancement of the Visual Response in the Monkey Striate Cortex and Frontal Eye Fields," *Journal of Neurophysiology* 39 (1976): 766–772.
- <sup>23</sup>Churchland, Ramachandran, and Sejnowski, "A Critique of Pure Vision," 51.
- <sup>24</sup>The quote is from Paul M. Churchland, *The Neurocomputational Perspective* (Cambridge, Mass.: MIT Press/Bradford Books, 1989). For an accessible introduction to artificial neural networks, see Paul Churchland, *The Engine of Reason, the Seat of the Soul* (Cambridge, Mass.: MIT Press, 1995).
- <sup>25</sup>This is known as a "subsumption" architecture, because the layers each constitute a complete behavior-producing system and interact only in simple ways such as by one layer subsuming (turning off) the activity of another or by one layer co-opting and hence "building in" the activity of another. See Brooks, "Intelligence Without Representation."
- <sup>26</sup>The retrieving is accomplished by a process of spreading activation among landmark encoding nodes. See Maja Mataric, "Navigating with a Rat Brain: A Neurobiologically Inspired Model for Robot Spatial Representation," in Jean-Arcady Meyer and Stewart W. Wilson, eds., *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior* (Cambridge, Mass.: MIT Press, 1991).

- <sup>27</sup>Clark, *Being There*, 47. The Mataric robot is based on actual rat neurobiology; see B. McNaughton and L. Nadel, "Hebb-Marr Networks and the Neurobiological Representation of Action in Space," in *Neuroscience and Connectionist Theory*, ed. Mark A. Gluck and David E. Rumelhart (Hillsdale, N.J.: Erlbaum Associates, 1990). Action-oriented representations bear some resemblance to what the ecological psychologist J. J. Gibson called "affordances." Yet Gibson himself would reject our emphasis on inner states and encodings, for an affordance (as he sees it) is the potential of use and activity that the local environment offers to a specific kind of being: chairs afford sitting (to humans), and so on. J. J. Gibson, *The Ecological Approach to Visual Perception* (Boston, Mass.: Houghton Mifflin, 1979). The philosopher Ruth Millikan has developed a nice account of action-oriented representation under the label "pushmi-pullyu representations"; see Ruth Millikan, "Pushmi-Pullyu Representations," in *Philosophical Perspectives 9: AI, Connectionism, and Philosophical Psychology*, ed. James E. Tomberlin (Atascadero, Calif.: Ridgeview Publishing, 1995).
- <sup>28</sup>Churchland, Ramachandran, and Sejnowski, "A Critique of Pure Vision," 60.
- <sup>29</sup>J. DiPelligrino, R. Klatzky, and B. McCloskey, "Time Course of Preshaping for Functional Responses to Objects," *Journal of Motor Behavior* 21 (1992): 307–316.
- <sup>30</sup>See G. Rizzolatti, L. Fadiga, and L. Fogassi, "Premotor Cortex and the Recognition of Motor Actions," *Cognitive Brain Research* 3 (1996): 131–141.
- <sup>31</sup>Marc Jeannerod, *The Cognitive Neuroscience of Action*, 191.
- <sup>32</sup>Donald Michie and Rory Johnson, *The Creative Computer: Machine Intelligence and Human Knowledge* (London: Penguin Books, 1984), 95.
- <sup>33</sup>See Pylyshyn, *The Robot's Dilemma*, for more discussion.
- <sup>34</sup>Dana H. Ballard et al., "Deictic Codes for the Embodiment of Cognition," *Behavioral and Brain Sciences* (forthcoming).
- <sup>35</sup>*Ibid.*, 6.
- <sup>36</sup>This example is borrowed from Andy Clark, "I am John's Brain," *Journal of Consciousness Studies* 2 (2) (1995): 144–148.
- <sup>37</sup>The story is detailed in Michael Triantafyllou and George Triantafyllou, "An Efficient Swimming Machine," *Scientific American* 272 (3) (1995): 64–71, and further discussed in Clark, *Being There*. A 49-inch, eight-segment anodized aluminum and lycra robot tuna is being used at MIT to test the details of the theory.
- <sup>38</sup>For more on cognitive extension, see Daniel C. Dennett, *Darwin's Dangerous Idea* (New York: Simon and Schuster, 1995), chaps. 12 and 13.
- <sup>39</sup>Dorothy Edwards, Carolyn Baum, and N. Morrow-Howell, "Home Environments of Inner-City Elderly: Do They Facilitate or Inhibit Function?" *The Gerontologist: Program Abstracts of the Forty-Seventh Annual Scientific Meeting* 34 (1) (1994): 64; Carolyn Baum, "Addressing the Needs of the Cognitively Impaired Elderly from a Family Policy Perspective," *American Journal of Occupational Therapy* 45 (7) (1991): 594–606.

- <sup>40</sup>Some of this additional structure is maintained and provided by family and friends. (But, similarly, much of our own wideware is provided by language, culture, and institutions that we do not ourselves create.)
- <sup>41</sup>See Clark, *Being There*, chart 10. Such counterbalancing is, as Marcel Kinsbourne has usefully reminded me, a somewhat delicate and complex matter. The mere provision of the various props and aids is useless unless the patient remains located in a stable, familiar environment. And the ability of different patients to make use of such added environmental structure itself varies according to the nature and extent of the neurally based deficit.
- <sup>42</sup>See Dawkins, *The Extended Phenotype*.
- <sup>43</sup>See W. Kutsch et al., "Wireless Transmission of Muscle Potentials During Free Flight of a Locust," *Journal of Experimental Biology* 185 (1993): 367–373. My thanks to Joe Faith for drawing this example to my attention.
- <sup>44</sup>See, for example, Esther Thelen and Linda B. Smith, *A Dynamic Systems Approach to the Development of Cognition and Action* (Cambridge, Mass.: MIT Press, 1994), Robert F. Port and Timothy Van Gelder, *Mind as Motion: Explorations in the Dynamics of Cognition* (Cambridge, Mass.: MIT Press/Bradford Books, 1995).
- <sup>45</sup>A move suggested by Edwin Hutchins, *Cognition in the Wild* (Cambridge, Mass.: MIT Press, 1995).
- <sup>46</sup>See Clark, *Being There*; Jim Crutchfield and Melanie Mitchell, "The Evolution of Emergent Computation," *Proceedings of the National Academy of Sciences* 92 (1995): 10742–10746; M. Mitchell, J. Crutchfield, and P. Hrabec, "Evolving Cellular Automata to Perform Computations," *Physica D* 75 (1994): 361–391; and Timothy Van Gelder and Robert F. Port, "It's About Time: An Overview of the Dynamical Approach to Cognition," in *Mind as Motion*, ed. Robert F. Port and Timothy Van Gelder (Cambridge, Mass.: MIT Press, 1995), 1–44.
- <sup>47</sup>Jerry A. Fodor, *The Language of Thought* (New York: Crowell, 1975). See also Jerry Fodor and Zenon Pylyshyn, "Connectionism and Cognitive Architecture: A Critical Analysis," *Cognition* 28 (1988): 3–71.
- <sup>48</sup>I am very grateful to Stephen Graubard and the participants at the *Dædalus* authors' meeting in Paris in October 1997 for a wealth of useful advice, good criticism, and wise counsel. Special thanks to Jean-Pierre Changeux, Marcel Kinsbourne, Vernon Mountcastle, Guilio Tonini, Steven Quartz, and Semir Zeki. As usual, any remaining errors are all my own.