

An embodied cognitive science?

Andy Clark

The last ten years have seen an increasing interest, within cognitive science, in issues concerning the physical body, the local environment, and the complex interplay between neural systems and the wider world in which they function. Yet many unanswered questions remain, and the shape of a genuinely physically embodied, environmentally embedded science of the mind is still unclear. In this article I will raise a number of critical questions concerning the nature and scope of this approach, drawing a distinction between two kinds of appeal to embodiment: (1) 'Simple' cases, in which bodily and environmental properties merely constrain accounts that retain the focus on inner organization and processing; and (2) More radical appeals, in which attention to bodily and environmental features is meant to transform both the subject matter and the theoretical framework of cognitive science.

Talk of embodiment and situatedness has become increasingly frequent in philosophy¹⁻⁵, psychology^{6,7}, neuroscience^{8,9}, robotics^{10,11}, education¹²⁻¹⁴, cognitive anthropology^{15,16}, linguistics^{4,17}, and in dynamical systems approaches to behavior and thought^{18,19}. There is clearly a shift in thinking but the nature and importance of the shift is surprisingly hard to pin down. What is 'embodied cognitive science' and how far can it take us? Let us look first at some of the varied roles that embodiment can play, as illustrated in the following examples.

Fish

Consider first the swimming ability of the Bluefin tuna. The Bluefin tuna is a swimming prodigy, but its aquatic capabilities – its ability to turn sharply, to accelerate quickly, and to reach such high speeds – have long puzzled biologists. Physically speaking, so it seemed, the fish should be too weak (by about a factor of seven) to achieve these feats. However, an explanation for this prodigious ability can be found in the use of embodied, environmentally embedded action by the tuna. Fluid dynamicists have suggested that the fish uses bodily action to manipulate and exploit the local environment (the water) so as to swim faster, accelerate more quickly, and so on²⁰. It appears that the tuna find and exploit naturally occurring currents so as to gain speed, and use tail flaps to create additional vortices and pressure gradients, which are then used for rapid acceleration and turning. The physical system whose functioning explains the prodigious swimming capacities of the Bluefin tuna is thus the fish-as-embedded-in, and as actively exploiting, its local environment.

Robots

Next consider a hopping robot, designed and built by Raibert and Hodgins in 1993 (Ref. 21). Their robots were designed to balance and move by hopping on a single leg – a pneumatic

cylinder with a kind of foot. To get the hopper to locomote – to move, balance and turn – involved solving a control problem that was radically impacted by the mechanical details, such as the elastic rebound that occurs when the leg hits the floor. The crucial control parameters included items such as the resting length of the leg spring, and degree of sideways tilt. To understand how the robot's 'brain' controls the robot's motions, a shift towards an embodied perspective is required. The controller must learn to exploit the rich intrinsic dynamics of the system. As Fred Keijzer recently put it, '...instead of thinking about [the] control system as a center for commands to be executed by actuators, the body and its movements are taken as a system with its own dynamic characteristics' (F. Keijzer, doctoral dissertation, University of Leiden, 1998). A similar shift in thinking can be applied to action routines in human infants and adults (see Box 1).

Vision

A further example can be found in research in animate or interactive vision^{22,23}. The key insight here is that the task of vision is *not* to build rich inner models of a surrounding 3-D reality²⁴, but rather to use visual information efficiently and cheaply in the service of real-world, real-time action. Researchers in animate and interactive vision thus reject what Churchland *et al.*²³ dub the paradigm of 'pure vision' – the idea (associated with work in classical AI and in the use of vision for planing) that vision is largely a means of creating a world model rich enough to let us 'throw the world away', allowing reason and thought to be focused upon the inner model instead. Real-world action, in these 'pure vision' paradigms, functions merely as a means of implementing solutions arrived at by pure cognition. The animate vision paradigm, by contrast, gives action a primary role²². Computational economy and temporal efficiency are

*A. Clark is at the
Philosophy/
Neuroscience/
Psychology (PNP)
Program, Washington
University in St Louis,
USA.*

tel: +1 314 935 7146
fax: +1 314 935 7349
e-mail: andy@
twinearth.wustl.edu

Box 1. The situated infant

Thelen and Smith have shown in some detail that infant walking abilities depend on complex interactions between neural states, the biomechanics of the legs, and local environmental parameters (Ref. a). This view is then contrasted with the vision of walking as the simple expression of a temporally staged set of genetically specified instructions (Ref. a, pp. 8–20, 263–266). Experimental support for the multi-factor view includes a demonstration of induced stepping in some ('non-stepping') seven-month-olds held upright upon a motorized treadmill. In this condition, the 'non-stepping' infants could even compensate for twin belts, each driving a leg at a different speed! The treadmill stepping, it was further discovered, depended crucially on the type of contact made between foot and belt: flat-foot contact induced stepping, whereas mere toe-contact failed. The explanation seems to be that the stepping behavior depends heavily on a spring-like biomechanical response (Ref. a, pp. 111–112). To uncoil the spring and propel the leg forward, it must first be stretched

fully back. Flat-foot contact with the moving treadmill creates this condition and initiates stepping, whereas toe contact yields insufficient back stretch. Infant stepping is thus a complex, multi-factor affair in which the target behavior 'emerges only when the central elements cooperate with the effectors – the muscles, joints, tendons – in the appropriate physical context' (Ref. a, p. 113).

Scott Kelso has demonstrated similar multi-factor profiles in adult motor routines such as the production of rhythmic finger motions (Ref. b.) For further commentary on both of these cases, see Clark (Ref. c.)

References

- a Thelen, E. and Smith, L. (1994) *A Dynamic Systems Approach to the Development of Cognition and Action*, MIT Press
- b Kelso, S. (1995) *Dynamic Patterns*, MIT Press
- c Clark, A. (1997) The dynamical challenge *Cognit. Sci.* 21, 461–481

purchased by a variety of tricks and ploys that exploit bodily action and local environment, including:

(1) the use of cheap, easy-to-detect environmental cues (e.g. searching for Kodak film in a drug store? Seek 'Kodak yellow');

(2) the use of active sensing (using motor action, guided by rough perceptual analysis, to seek further inputs yielding *better* perceptual data. Move head and eyes for better depth perception, etc.);

(3) the use of repeated consultations of the outside world in place of rich, detailed inner models.

This approach suggests that vision is a highly active and intelligent process. It is not the passive creation of a rich inner model, but rather the active retrieval (typically by moving the high resolution fovea to sequential locations in the visual scene) of useful information *as it is needed* from the constantly present real-world scene. Ballard *et al.*²⁵ speak of 'just-in-time representation', while the roboticist, Rodney Brooks, has coined the slogan 'The world is its own best model'¹¹. Biological vision thus gears its computational activity closely and sparingly to the task at hand, making the most efficient use of the persisting external scene.

Action and affordance

Related insights stem from the work of J.J. Gibson^{26,27} and the ecological psychology movement^{7,28,29}. This approach stresses bodily movement, ecological context and the action-relevant information available in the perceptual array. A central organizing construct is the concept of an 'affordance'²⁶. Affordances are the possibilities for use, intervention and action which the physical world offers a given agent and are determined by the 'fit' between the agent's physical structure, capacities and skills and the action-related properties of the environment itself³⁰.

A simple but illustrative example is Lee and Reddish's account of how diving birds, such as plovers and gannets, are able to close their wings at exactly the right moment before hitting the surface of the water in pursuit of a fish³¹. Such behavior is possible because there is available in the optic array, a higher-order invariant that allows the control of such action. This quantity (which involves the relative rate of expansion of the

image in the optic array) accurately predicts time-to-impact and can be used to time wing closure. Other behaviors, such as the timing of an athlete's jump or a stroke in tennis, make use of a similar quantity (see Ref. 32 for a review).

A similar approach can help explain how an outfielder in baseball positions him/herself to catch a fly ball. It used to be thought that this problem required complex calculations of the arc, acceleration and distance of the ball. More recent work, however, suggests a computationally simpler strategy³³. Put simply, the fielder continually adjusts his or her run so that the ball never seems to curve towards the ground, but instead appears to move in a straight line in his or her visual field. By maintaining this strategy, the fielder should be guaranteed to arrive in the right place at the right time to catch the ball.

Notice the difference between these two models. In the traditional model, the brain takes in data, performs a complex computation that solves the problem (where will the ball land?) and then instructs the body where to go. This is a linear processing cycle: perceive, compute and act. In the second model, the problem is not solved ahead of time. Instead, the task is to maintain, by multiple, real-time adjustments to the run, a kind of *co-ordination* between the inner and the outer worlds. Such co-ordination dynamics constitute something of a challenge to traditional ideas about perception and action: they replace the notion of rich internal representations and computations, with the notion of less expensive strategies whose task is not first to represent the world and then reason on the basis of the representation, but instead to maintain a kind of adaptively potent equilibrium that couples the agent and the world together. Whether such strategies are genuinely non-representational and non-computational, or suggestive of different *kinds* of representation ('action-oriented representations') and more efficient forms of computation, is a difficult question whose resolution remains uncertain^{2,6,7,18,34–36}.

What is clear, however, is that tuning to higher-order invariants can help explain a wide variety of adaptive responses, including visually guided locomotion^{37,38}, rhythmic movement^{18,39}, and the capacity to grasp and wield objects⁷ (e.g. hammers, golf clubs). In all these cases behavioral success

involves locking on to simple (but often far from obvious) properties of the environment made available in the perceptual array (see also Box 1).

Beyond adaptive coupling?

The full implications and significance of these embodied and embedded approaches remain to be determined, and there are a number of difficulties that clearly remain to be resolved. An immediate question is, to what extent, if at all, can the embodied, embedded approach contribute to our understanding of so-called ‘representation-hungry’ problem-solving⁴⁰? To illustrate this idea, consider the much simpler notion of ‘adaptive coupling’. Adaptive coupling occurs when a system (typically a plant or animal) evolves a mechanism that allows it to track the behavior of another system (a predator, or a source of food or energy).

To borrow an example from Brian Cantwell Smith, the sunflower has evolved to track the daily motion of the sun across the sky⁴¹. The sunflower thus has states that co-vary with solar position, and this is what they are *meant* (evolutionarily speaking) to do. But does the sunflower thereby exhibit cognition or mentality or intentionality; does it have internal representations? A common (and I think correct) intuition is that it does not. There is nothing cognitive occurring. One reason for thinking this is that cognition has been taken to involve the capacity to relate to an ‘intentional object’⁴² – and this means, in part, an object that might not be present-at-hand or that might not even exist. The sunflower, by contrast, tracks the sun only when the sun is (in a certain sense, at least) actually there. More precisely, the sunflower exhibits its behavior when there is an ongoing external physical trace to which it can adaptively couple.

The mark of the cognitive, then, is the capacity to engage in something like off-line reason⁴³ – reasoning in the absence of that which our thoughts concern. Classical (‘disembodied’) cognitive science accounted well for this, by positing an inner realm richly populated with internal tokens that *stood for* external objects and states of affairs. Thus, it was able to offer a simple account of behavioral co-ordination in the absence of any external physical trace or perceptually available higher-order invariant (for some excellent discussion of the issues confronting a Gibsonian approach to cognition, see Kirsh⁴⁴ and Van Leeuwen⁴⁵).

One promising advance is the suggestion that embodied cognitive science might treat off-line reason as something like *simulated* sensing and acting, thus preserving the special flavor of embodied problem-solving alongside a high degree of ability to decouple from the environment. The most developed version of this strategy is probably the mobile robotics work of Lynn Andrea Stein at the MIT Artificial Intelligence lab⁴⁶. Stein uses as her platform a mobile robot (designed and implemented by Maja Mataric⁴⁷) named TOTO. TOTO uses ultrasonic range sensors to detect walls, corridors and other obstacles and is able to use its physical explorations to build an inner map of its environment, which it can then use to revisit previously encountered locations on command. TOTO’s internal ‘map’ is, however, rather special in that it encodes geographic information in an ‘action-oriented’ way (Ref. 2, p. 49), combining information about the robot’s movement and correlated perceptual input. TOTO’s inner

mechanisms thus record landmarks as a combination of robotic motion and sonar readings, so that a corridor might be encoded as a combination of forward motion and a sequence of short, lateral sonar distance readings. The stored ‘map’ is thus perfectly formatted to act as a direct controller of embodied action: using the map to find a route and generating a plan of actual robot movements is therefore a single computational task.

TOTO is adept at interacting with the local environment, and can even, in a weak sense, ‘track’ that which is not present-to-hand: it can return, on command, to a previously encountered location. TOTO cannot, however, be prompted to track or ‘think about’ any location that it has not previously visited. METATOTO⁴⁶ builds on the original TOTO architecture to create a system capable of finding its way, on command, to locations that it has never previously encountered. It does so by using the TOTO architecture off-line, so as to support the exploration, in ‘imagination’, of a totally virtual ‘environment’. When METATOTO is ‘imagining’, it deploys exactly the same machinery that (in TOTO, and in METATOTO on-line) normally supports physical interactions with the real world. The difference lies solely at the lowest-level interface: where TOTO uses sonar to act and navigate in a real world, METATOTO uses simulated sonar to explore a virtual world (including a virtual robot body). Stein then goes on to imagine linguistic directions interfacing with this virtual realm by translating descriptions such as ‘the second left’ into TOTO (METATOTO)-style action-based encodings, such as ‘short sonar left, long sonar left, short sonar left, long sonar left’ (Ref. 46, p. 404).

METATOTO uses the basic behavior-producing architecture of TOTO, but includes a program that can take, for example, a floor plan or map and use it to stimulate the robot’s sensors in the way they would be stimulated if the robot were actually locomoting along a given route on the map. The map can thus induce sequences of experiences that are qualitatively similar to those generated by real sensing and acting, and this allows METATOTO to profit from ‘virtual experiences’, just as TOTO profits from real experience. Once the sensors and motors are restored to real world input and action, on-line METATOTO can immediately find its way to a target location it has not actually (but merely virtually) visited.

We should now ask two, related questions. How different is this account from more traditional solutions? And will it work for all kinds of off-line reasoning or only some? The first question, it seems to me, leads to a mild dilemma. For the simulation-based account looks most clearly different (from traditional accounts involving inner-world models) only insofar as it treats planning as, quite literally, imagined interaction. Thus, Stein notes that ‘While traditional planners use an abstracted world and plan operators distinct from the actual robot controls, our system uses the robotic architecture itself’ (Ref. 46, p. 396). To support this claim, Stein reminds us that METATOTO works by simulating both sensors and actuators, and that simulation runs create the kinds of feedback (short and long sonar signals, etc.) that would be received from the actual world, were the robot actually to change position. There are, of course, some idealizations: the simulated motion is, for example, straight and precise, unaffected by

Box 2. Dynamical systems theory

The dynamical systems approach has been gaining ground in cognitive scientific treatments of cognition and adaptive behavior (Refs a–d). The dynamical approach focuses on the evolution of a system over time, and is particularly well-suited to dealing with cases in which a system or component *a* is constantly affecting and being affected by another system or component *b* (which might likewise be continuously sensitive to item *c* and so on). An example might be the process of returning a tennis serve: the location of the ball and the other player (and perhaps your partner, in doubles) are constantly changing, and simultaneously, you are moving and acting, which is affecting the other players, and so on. Put simply, ‘everything is simultaneously affecting everything else’ (Ref. a, p. 23). Such densely coupled events, insofar as they prove characteristic of all or some human cognitive activity, seem ill-understood in traditional cognitive scientific terms. By contrast, the apparatus of dynamical systems theory, with its notions of state spaces, coupled systems, trajectories in state space, collective variables and more is expressly designed to deal with such simultaneous interactive complexity (Ref. e). It provides a set of mathematical and conceptual tools that support a *geometric* understanding of the space of possible total system behaviors. Such analyses have proven useful in understanding the activity of simple robots (Ref. f), infants (Ref. d), and adults (Ref. c).

But what mileage can we get from such analyses once we leave the domain of on-the-spot adaptive coupling and turn to various forms of off-line reason and cogitation? One exciting possibility,

recognized by van Gelder and Port (Ref. b) and pursued by Melanie Mitchell and her colleagues (Ref. g), is that the dynamical approaches might transform and enrich (rather than displace) our computational and representational understandings, perhaps by identifying complex, temporally extended dynamical patterns (chaotic attractors, limit cycles, values of collective variables, etc.) as the vehicles of specific representational contents and as the implementation of so-called ‘emergent’ computational processes (Ref. h).

References

- a Port, R. and van Gelder, T., eds (1995) *Mind as Motion: Dynamics, Behavior, and Cognition*, MIT Press
- b van Gelder, T. and Port, R. (1995) It’s about time: an overview of the dynamical approach to cognition, in *Mind as Motion: Explorations in the Dynamics of Cognition* (Port, R. and van Gelder, T., eds), pp. 1–44, MIT Press
- c Kelso, S. (1995) *Dynamic Patterns*, MIT Press
- d Thelen, E. and Smith, L. (1994) *A Dynamic Systems Approach to the Development of Cognition and Action*, MIT Press
- e Abraham, R. and Shaw, C. (1992) *Dynamics: The Geometry of Behavior*, Addison–Wesley
- f Beer, R. (1995) A dynamical systems perspective on environment agent interactions *Artif. Intell.* 72, 173–215
- g Mitchell, M., Crutchfield, J. and Hrabner, P. (1994) Evolving cellular automata to perform computations *Physica D* 75, 361–391
- h Crutchfield, J. and Mitchell, M. (1995) The evolution of emergent computation *Proc. Natl. Acad. Sci. U. S. A.* 92, 10742–10746

the dinks and slopes found in the real world. But, overall, METATOTO does indeed rely on the simulation of sensorimotor experience rather than on abstract kinds of reasoning and planning.

What remains unclear, however, is the *scope* of this kind of solution. The reader might, for example, try the following exercise in abstract off-line reasoning: consider whether US gun manufacturers should be held liable for having knowingly manufactured more guns than the legal market could possibly account for? Here it is not clear how rich sensorimotor simulation could possibly account for all the kinds of moral and abstract reasoning required – reasoning about rights, implications, responsibilities, economics, and so on. Nor is this just a point about how things seem introspectively. Rather, it is hard to see how sensorimotor simulation could in principle account for all the kinds of thought and reasoning that the problem demands. Simulated acting and sensing may well play a role – and perhaps even an essential role⁴ – in our reasoning. But the capacity to examine arguments, to judge what follows from what, and to couch the issues in the highly abstract terms of a fundamental moral debate (using concepts like ‘liability’, ‘reasonable expectation’, ‘acceptable risk’, etc.) does not obviously lend itself to an analysis in terms of literal sensorimotor simulation as in METATOTO. Perhaps there are other, less direct ways to depict high-level cognition as dependent on simulated sensing and acting⁴⁸. But (and this is the mild dilemma) it does seem that the more decoupled and abstract the target contents become, *either* the less applicable the sensorimotor simulation strategy is, *or* the less clearly it can then be differentiated from the more traditional approaches it seeks to displace.

Simple versus radical embodiment

In addition to asking how far the embodied approach can go, we should also ask to what extent it is a genuinely radical alternative to more traditional views. To focus this concern, I would like to distinguish two different ways to appeal to facts about embodiment and environmental embedding. The first, which I will call ‘simple embodiment’, treats such facts as, primarily, *constraints upon a theory of inner organization and processing*. The second, which I will call ‘radical embodiment’ goes much further and treats such facts as *profoundly altering the subject matter and theoretical framework of cognitive science*. The distinction between the simple and radical forms is, however, not absolute, and many (perhaps most) good research programs end up containing elements of both.

Examples of simple embodiment abound in the literature. A good deal of work in interactive vision, for example, still relies heavily on internal representations, computational transformations, and abstract data structures^{22,25}. There is much talk for example, of ‘inner databases’, of ‘internal featural representations’ (of color, shape and so on), of ‘high-dimensional feature vectors’, and so on. Attention to the roles of body, world and action, in such cases, is merely a methodological tool aimed at getting the *internal* data-structures and operations right. Churchland *et al.*’s vision of a ‘motocentric’ rather than ‘visuocentric’ cognitive science has, I suspect, a similar goal (Ref. 23, p. 60). Maja Mataric’s⁴⁷ work on TOTO has this flavor, insofar as it concentrates attention on an inner representational resource (the map/controller) and is exploring the ways in which usefulness in the guidance of real-world action can both constrain and inform the nature of inner representations and processing. The same

applies to the recent important work on the role of bodily metaphors⁴ in abstract, high-level cognition: here, too, the goal is to give an account of the inner representational realm, but one informed by the evolutionary and developmental roles of bodily experience.

The source of much recent excitement, however, are the striking claims involving ‘radical embodiment’. Such claims can be found in work by Tim Van Gelder⁴⁹, Thelen and Smith⁶, Kelso¹⁸, Varela *et al.*¹, Turvey and Carello⁷ and others. These accounts of radical embodiment all involve one or more of the following claims:

(I) that understanding the complex interplay of brain, body and world requires new analytic tools and methods, such as those of dynamical systems theory^{6,18,49} (see Box 2)

(II) that traditional notions of internal representation and computation are inadequate and unnecessary^{6,7,49}

(III) that the typical decomposition of the cognitive system into a variety of inner neural or functional subsystems is often misleading, and blinds us to the possibility of alternative, and more explanatory, decompositions that cut across the traditional brain–body–world divisions^{3,5,6,15}.

Closely related to these three points is the idea that even the subject matter of cognitive science needs to be re-thought. A mature science of the mind, it now seems, targets not (or not only) the individual, inner organization of intelligence but the bodily and environmentally extended organizations responsible for adaptive success^{2,12,15}.

Some support for claims (I) and (II) may be found in the work on infant motor development^{30,37} (Box 1), adult motor actions^{7,18}, and mobile robotics^{9–11,21}. The support is weak, however, because the solutions that appear most non-computational, representation-free, and open to dynamical analysis (see Box 2) usually involve cases of adaptive coupling, and do not directly confront ‘representation-hungry’ problems. Here, we must simply suspend judgement and await empirical advances.

My own guess, however, is that as tasks become more representation-hungry – more concerned with the distal, abstract and non-existent – we will see more and more evidence of *some kinds* of internal representation and inner models. It is at exactly this point that the possibility of a middle ground between simple and radical versions of embodiment becomes apparent. For these new kinds of internal representation might differ from familiar forms both in their contents (being more ‘deictic’²⁵ and action-oriented²) and in the nature of their inner ‘vehicles’ (perhaps using temporally extended processes and complex dynamical regularities as inner ‘tokens’^{2,50,51}).

A good prospect for the supporters of radical embodiment could, however, lie in claim (III) (that there might be alternative systemic decompositions). An example – which also demonstrates how a single research program can combine elements of simple and radical embodiment – is Ballard *et al.*'s²⁵ use of a notion of ‘deictic pointers’. A pointer, in artificial intelligence, is an inner state, which can act both as an object of computation and as a ‘key’ for retrieving additional data-structures or information. Deictic pointers, as Ballard *et al.* describe them, are physical actions – such as foveating a certain location in visual space – that play a similar kind of functional role. The very act of foveating, it is suggested, may be used to temporarily bind color to location, or to direct

a reaching motion to a target. In such cases, Ballard *et al.* suggest, ‘the external world is analogous to computer memory’ and ‘changing gaze is analogous to changing the memory reference in a silicon computer’ (Ref. 28, p. 725). The computational organization relevant to cognition is here depicted as literally spread across neural, bodily and environmental elements.

In thinking about ‘higher’ cognition and advanced human reason, it might likewise prove fruitful to consider the literal extension of the cognitive system to include aspects of the local environment. In this vein, Clark^{2,52} and Hutchins¹⁵, following Vygotsky⁵³, Bruner⁵⁴, Dennett⁵⁵ and others, have argued that just as basic forms of real-world success turn on the interplay between neural, bodily and environmental factors, so advanced cognition turns – in crucial respects – upon the complex interplay between individual reason, artifact and culture. The simplest illustration of this idea is probably the use of artifacts such as pen and paper to support or ‘scaffold’ human performance^{34,56–58}. Most of us, armed with pen and paper, can, for example, solve multiplication problems that would baffle our unaided brains. In so doing we create external symbols (numerical inscriptions) and use external storage and manipulation so as to reduce the complex problem to a sequence of simpler, pattern-completing steps that we already command⁵⁸. On this model, then, it is the combination of our biological computational profile with the fundamentally different properties of a structured, symbolic, external resource that is a key source of our peculiar brand of cognitive success^{53,55}. The external environment, actively structured by us, becomes a source of cognition-enhancing ‘wideware’⁵² – external items (devices, media, notations) that scaffold and *complement* (but usually do not replicate) biological modes of computation and processing, creating extended cognitive systems whose computational profiles are quite different from those of the isolated brain. Hutchins, for example, provides a lucid and detailed account of the way multiple biological brains, tools (such as sextants and alidades), and media (such as maps and charts) combine to make possible the act of ship navigation¹⁵. In Hutchins’ words, such tools and media ‘permit the users to do the tasks that need to be done while doing the kinds of things people are good at: recognizing patterns, modeling simple dynamics of the world, and manipulating objects in the environment’ (Ref. 15, p. 155).

In short, the world of artifacts, texts, media, and even cultural practices and institutions⁵⁹, might be for us what the actively created whirls and vortices are for the Bluefin tuna. Human brains, raised in this sea of cultural tools⁵⁵ might develop strategies for advanced problem solving that ‘factor in’ these external resources as profoundly and deeply as the bodily motions of the tuna factor in and maximally exploit the reliable properties of the surrounding water.

Recognizing the complex ways in which human thought and reason exploit the presence of external symbols and problem-solving resources, and unraveling the ways in which biological brains couple themselves with these very special kinds of ecological objects, is surely one of the most exciting tasks confronting the science of embodied cognition – and one that might shed great light on the role of embodiment in more abstract cognitive domains.

Outstanding questions

- Is cognition truly seamless – implying a gentle, incremental trajectory linking fully embodied responsiveness to abstract thought and off-line reason? Or is it a patchwork quilt, with jumps and discontinuities and with very different *kinds* of processing and representation serving different needs?
- What role does public language play in the transition from simple adaptive coupling to heavy-duty cognition?
- Insofar as we depend heavily on cultural artifacts (pen, paper, PC) to augment and enhance biological cognition, what should we say about their *origins*? If we do indeed 'make our world smart so that our brains can be dumb in peace'², just how *did* dumb brains create such a smart world?
- Are the tools of dynamical systems theory replacements for, or merely additions to, the familiar arsenal of inner models, maps, representations and computations?
- If we follow the embodied, embedded approach to its natural conclusions, do we lose sight of the differences between perception, reason and action? If not, just how do we reconstruct them? Do we begin to lose sight of the distinction between agents and the worlds in which they think and act?

Conclusions

Embodied, environmentally embedded approaches have a lot to offer cognitive science. It is increasingly clear that, in a wide variety of cases, the individual brain should not be the sole locus of cognitive scientific interest. Cognition is not a phenomenon that can be successfully studied while marginalizing the roles of body, world and action.

The major challenge for the vision of 'radical embodiment' described here lies with the class of 'representation-hungry' problems and the phenomena of off-line, abstract, and environmentally-decoupled reason. It is important *not* to conclude, however, that facts about embodiment impact only our ideas about low-level sensorimotor processes. In the human case, at least, we seem to find at all levels a mixture of highly 'embodied, embedded' strategies and apparently much more abstract and potentially de-coupled strategies, with the creation and manipulation of external symbolic items often functioning as a kind of bridge between the two. It thus seems likely that one key to understanding the nature and potency of human thought and reason lies precisely in understanding the complex relations and interactions between these various types of strategy and resource⁴¹. (Human language skills – a topic I have deliberately avoided in this review – are a case in point; words and text are both real, external objects that we can encounter and manipulate *and* key instruments of inner, abstract, environmentally decoupled reason^{55,60}.)

The gulf between the embodied, embedded skills of the Bluefin tuna and the more de-coupled skills of the moralists and mathematicians remains. But the size, nature and significance of this gap are matters for further research and debate. At the very least, an embodied cognitive science must now look beyond the on-line production of tuned motor responses to the creation, maintenance and transformation of the inner and outer states that together allow us to know the world as an arena for embodied action.

Acknowledgements

Thanks to Tim Van Gelder, Esther Thelen, Dan Dennett, Randy Beer, Michael Wheeler, Lynn Stein, Brian Cantwell Smith, Dominic Murphy, Timo Jarvilehto and the anonymous referees for invaluable help and suggestions on these topics.

References

- 1 Varela, F. et al. (1991) *The Embodied Mind*, MIT Press
- 2 Clark, A. (1997) *Being There: Putting Brain, Body and World Together Again*, MIT Press
- 3 Hurley, S. (1998) *Consciousness in Action*, Harvard University Press
- 4 Lakoff, G. and Johnson, M. (1999) *Philosophy in the Flesh*, MIT Press
- 5 Haugeland, J. (1998) Mind embodied and embedded, in *Having Thought* (Haugeland, J., ed.), pp. 3–38, MIT Press
- 6 Thelen, E. and Smith, L. (1994) *A Dynamic Systems Approach to the Development of Cognition and Action*, MIT Press
- 7 Turvey, M. and Carello, C. (1995) Some dynamical themes in perception and action, in *Mind as Motion* (Port, R. and Van Gelder, T., eds), pp. 373–401, MIT Press
- 8 Damasio, A. (1994) *Descartes' Error*, Grosset/Macmillan
- 9 Chiel, H. and Beer, R. (1997) The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment *Trends Neurosci.* 20, 553–557
- 10 Arkin, R. (1998) *Behavior-Based Robotics*, MIT Press
- 11 Brooks, R. (1991) Intelligence without representation *Artif. Intell.* 47, 139–159
- 12 Resnick, M. (1994) *Turtles, Termites and Traffic Jams*, MIT Press
- 13 Anderson, J., Reder, L. and Simon, H. (1996) Situated learning and education *Educational Researcher* 25, 5–11
- 14 Greeno, J. (1996) On claims that answer the wrong questions (Response to Anderson et al.) *Educational Researcher* 25, 5–11
- 15 Hutchins, E. (1995) *Cognition in the Wild*, MIT Press
- 16 Suchman, A. (1987) *Plans and Situated Actions*, Cambridge University Press
- 17 Regier, T. (1996) *The Human Semantic Potential*, MIT Press
- 18 Kelso, S. (1995) *Dynamic Patterns*, MIT Press
- 19 Port, R. and van Gelder, T., eds (1995) *Mind as Motion: Dynamics, Behavior, and Cognition*, MIT Press
- 20 Triantafyllou, M. and Triantafyllou, G. (1995) An efficient swimming machine *Sci. Am.* 272, 64–71
- 21 Raibert, M. and Hodgins, T. (1993) Legged robots, in *Biological Neural Networks in Invertebrate Neuroethology and Robotics* (Beer, R.D., Ritzmann, R.E. and McKenna, T., eds), pp. 319–354, Academic Press
- 22 Ballard, D. (1991) Animate vision *Artif. Intell.* 48, 57–86
- 23 Churchland, P. et al. (1994) A critique of pure vision, in *Large-Scale Neuronal Theories of the Brain* (Koch, C. and Davis, J., eds), MIT Press
- 24 Marr, D. (1982) *Vision*, W.H. Freeman
- 25 Ballard, D. et al. (1997) Deictic codes for the embodiment of cognition *Behav. Brain Sci.* 20, 723–767
- 26 Gibson, J.J. (1996) *The Senses Considered as Perceptual Systems*, Houghton-Mifflin
- 27 Gibson, J.J. (1979) *The Ecological Approach to Visual Perception*, Houghton-Mifflin
- 28 Gibson, E.J. (1982) The concert of affordances in development, in *Minnesota Symposium on Child Psychology* (Vol. 15) (Collins, W., ed.), pp. 55–81, Huirsdale
- 29 Michaels, C. and Cavello, C. (1981) *Direct Perception*, Prentice-Hall
- 30 Adolph, K. (1995) Psychophysical assessment of toddlers' ability to cope with slopes *J. Exp. Psychol. Hum. Percept. Perform.* 21, 734–750
- 31 Lee, D. and Reddish, P. (1981) Plummeting gannets: a paradigm of ecological optics *Nature* 293, 293–294
- 32 Tresilian, J.R. (1999) Visually timed action: time-out for tau? *Trends Cognit. Sci.* 3, 301–310
- 33 McBeath, M., Shaffer, D. and Kaiser, M. (1995) How baseball outfielders determine where to run to catch fly balls *Science* 268, 569–573
- 34 Clark A. (1995) Moving minds: situating content in the service of real-time success, in *Philosophical Perspectives* (Vol. 9: *A-I, Connectionism and Philosophical Psychology*) (Tomberlin, J., ed.), pp. 89–104, Ridgeview
- 35 Rutkowska, J. (1993) *The Computational Infant*, Harvester Press
- 36 Neisser, U. (1976) *Cognition and Reality*, W.H. Freeman
- 37 Adolph, K. and Eppler, M. (1998) Development of visually guided locomotion *Ecol. Psychol.* 10, 303–321
- 38 Adolph, K., Eppler, M. and Gibson, E. (1993) Crawling versus walking infants' perception of affordance for locomotion over sloping surfaces *Child Dev.* 64, 1158–1174
- 39 Hatsopoulos, N. and Warren, W. (1996) Resonance tuning in rhythmic arm movements *J. Mot. Behav.* 78, 3–14

- 40 Clark, A. and Toribio, J. (1995) Doing without representing? *Synthese* 101, 401–431
- 41 Smith, B.C. (1996) *On the Origin of Objects*, MIT Press
- 42 Brentano, F. (1874) *Psychologie vom Empirischen Standpunkt* (transl. 1973, *Psychology from an Empirical Standpoint*), Taylor & Francis, London
- 43 Clark, A. and Grush, R. Towards a cognitive robotics *Adaptive Behav.* (in press)
- 44 Kirsh, D. (1992) Today the earwig, tomorrow man?, in *Foundations of Artificial Intelligence* (Kirsh, D., ed.), pp. 161–184, MIT Press
- 45 Van Leeuwen, C. (1998) Perception, in *A Companion to Cognitive Science* (Bechtel, W. and Graham, G., eds), pp. 265–281, Blackwell
- 46 Stein, L. (1994) Imagination and situated cognition *J. Exp. Theoret. Artif. Intell.* 6, 393–407
- 47 Mataric, M. (1992) Integration of representation into goal-driven behavior based robots *IEEE J. Robotics Automation* 8, 304–312
- 48 Barsalou, L. Perceptual symbol systems *Behav. Brain Sci.* (in press)
- 49 van Gelder, T. (1995) What might cognition be, if not computation? *J. Philos.* XCII, 345–381
- 50 Clark, A. (1997) The dynamical challenge *Cognit. Sci.* 21, 461–481
- 51 Mitchell, M., Crutchfield, J. and Hrabec, P. (1994) Evolving cellular automata to perform computations *Physica D* 75, 361–391
- 52 Clark, A. (1998) Where brain, body and world collide *Daedalus* 127, 257–280
- 53 Vygotsky, L. (1986) *Thought and Language*, MIT Press
- 54 Bruner, T. (1990) *Acts of Meaning*, Harvard University Press
- 55 Dennett, D. (1996) *Kinds of Minds*, Basic Books
- 56 Wood, D., Bruner, J. and Russ, G. (1976) The role of tutoring in problem solving *J. Child Psychol. Psychiatry* 17, 89–100
- 57 Norman, D.A. (1993) *Things That Make Us Smart*, Addison-Wesley
- 58 McClelland, J. et al. (1986) The appeal of parallel distributed processing, in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (Vol. II) (McClelland, J., Rumelhart, D. and PDP Research Group, eds), pp. 3–44, MIT Press
- 59 Clark, A. (1997) Economic reason: the interplay of individual learning and external structure, in *The Frontiers of the New Institutional Economics* (Drobak, J., ed.), Academic Press
- 60 Clark, A. (1998) Magic words: how language augments human computation, in *Language and Thought* (Boucher, J. and Carruthers, P., eds), pp. 162–183, Cambridge University Press

Neuromodulation: acetylcholine and memory consolidation

Michael E. Hasselmo

Clinical and experimental evidence suggests that hippocampal damage causes more severe disruption of episodic memories if those memories were encoded in the recent rather than the more distant past. This decrease in sensitivity to damage over time might reflect the formation of multiple traces within the hippocampus itself, or the formation of additional associative links in entorhinal and association cortices. Physiological evidence also supports a two-stage model of the encoding process in which the initial encoding occurs during active waking and deeper consolidation occurs via the formation of additional memory traces during quiet waking or slow-wave sleep. In this article I will describe the changes in cholinergic tone within the hippocampus in different stages of the sleep–wake cycle and will propose that these changes modulate different stages of memory formation. In particular, I will suggest that the high levels of acetylcholine that are present during active waking might set the appropriate dynamics for encoding new information in the hippocampus, by partially suppressing excitatory feedback connections and so facilitating encoding without interference from previously stored information. By contrast, the lower levels of acetylcholine that are present during quiet waking and slow-wave sleep might release this suppression and thereby allow a stronger spread of activity within the hippocampus itself and from the hippocampus to the entorhinal cortex, thus facilitating the process of consolidation of separate memory traces.

This article will describe some of the physiological and neurochemical mechanisms that might mediate episodic memory consolidation. Firstly, I will review the neuropsychological and computational evidence for a two-stage model of

memory consolidation. I will then describe electrophysiological data that support this concept and suggest that the two stages are linked to different stages of the sleep–wake cycle when encoding and consolidation might occur. In the final section, I

Michael E. Hasselmo
is at the Department
of Psychology,
Boston University,
64 Cummington
Street, Boston,
MA 02215, USA.

tel: +1 617 495 3875
fax: +1 617 495 3728
e-mail: hasselmo@bu.
edu