A shortened and amended version appears as "An Embodied Cognitive Science?" TRENDS IN COGNITIVE SCIENCES 3:9:1999 p. 345-351

Embodiment: From Fish to Fantasy

Andy Clark

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

Abstract

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. "Physically embodied, environmentally embedded" approaches thus loom large on the contemporary cognitive scientific scene. Yet many unanswered questions remain, and the shape of a genuinely embodied, embedded science of the mind is still unclear. I begin by sketching a few examples of the approach, and then raise a variety of critical questions concerning its nature and scope. A distinction is drawn between two kinds of appeal to embodiment: 'simple' cases, in which bodily and environmental properties merely constrain accounts that retain the focus on inner organization and processing, and 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.

Things of the Flesh

Cognitive science is bringing the body out of the closet. Talk of embodiment, and situatedness looms increasingly large in philosophy[1-5], psychology[6-7], neuroscience[8,9], robotics[10-11], education[12-14], cognitive anthropology[15,16], linguistics[4,17], and in dynamical systems approaches to behavior and thought[18,19]. Something is afoot. But it is surprisingly – and frustratingly – hard to get a firm grip on exactly what it is. What is "embodied cognitive science" and how far can it take us?

We can begin by illustrating some of the varied roles that embodiment can play.

Fish

Consider first the swimming know-how of the Bluefin Tuna. The Bluefin Tuna is a swimming prodigy, but its aquatic capacities – its ability to turn sharply, to move off quickly, and to reach such high speeds – have puzzled biologists. Physically speaking, so

it seemed, the fish should be too weak (by about a factor of 7) to achieve these feats. But the explanation is not magic, but the canny use of embodied, environmentally embedded action. Fluid dynamicists at MIT[20] suggest that the fish use bodily action to manipulate and exploit the local environment (water) so as to swim faster, blast off more quickly, and so on. They (the fish) 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 quick take-off and rapid turns. The physical system whose functioning explains the prodigious swimming capacities is thus the fish-as-embedded-in, and as actively exploiting, its local environment.

Robots

Next in line is a hopping robot. Raibert and Hodgins, back in 1993[21], designed and built robots that balance and move by hopping on a single leg – basically, a pneumatic cylinder with a kind of foot. To get the hopper to locomote – to move, balance, and turn – involves solving a control problem that is radically impacted by the mechanical details, such as the elastic rebound when the leg hits the floor. The crucial control parameters include items such as leg spring rest length, and degree of sideways tilt. To understand how the robot's 'brain' controls the robot's motions involves a shift towards an embodied perspective. The controller must learn to exploit the rich intrinsic dynamics of the system. As Fred Keijzer (in his doctoral dissertation, "The Generation of Behavior", University of Leiden, 1998, p. 125) 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." A similar story is told concerning action routines in human infants and adults (see box 1).

Vision

Or consider vision. There is now a growing body of work devoted to so-called Animate[22] (or interactive[23]) vision. The key insight here is that the task of vision is *not* to build rich inner models of a surrounding 3D reality[24], but rather to use visual information efficiently and cheaply in the service of real-world, real-time action. Animate and Interactive vision thus rejects what Churchland, Ramachandran, et al[23] nicely dub the paradigm of "pure vision" – the idea (associated with work in classical AI and in the use of vision for planning), that vision is largely a means of creating a world model rich enough to let us "throw the world away', targeting reason and thought 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[22], by contrast, gives action a starring role. Here, computational economy and temporal efficiency is purchased by a variety of bodily action and local environment exploiting tricks and ploys including

- the use of cheap, easy-to-detect (possibly idiosyncratic) environmental cues (searching for Kodak film in a drug store? Seek 'Kodak yellow').
- the use of active sensing (use motor action, guided by rough perceptual analysis, to seek further inputs yielding *better* perceptual data. Move head and eyes for better depth perception, etc.).
- the use of repeated consultations of the world in place of rich, detailed inner models.

Vision, this body of work suggests, is a highly active and intelligent process. It is not the passive creation of a rich inner model, so much as the active retrieval (typically by moving the high resolution fovea in a saccade) 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]. The moral is clear: vision makes the most of the persisting external scene, and gears its computational activity closely and sparingly to the task at hand.

Action and Affordance

Related insights stem from the work of J. J. Gibson[26,27] and the ecological psychology movement[7,28,29]. This work 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 itself[30].

A simple but illustrative example is Lee and Reddish's[31] 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 fish. Such behavior is possible because there is available in the optic array, a higher order invariant which is ideal for 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 closing, as well as other behaviors, such as the timing of an athlete's jump.

A similar story can help explain how an outfielder positions 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 strategy[32]. Roughly, you continually adjust your run so that the ball never seems to curve towards the ground, but instead looks to move in a straight line

in your visual field. By maintaining such a co-ordination, it is mathematically guaranteed that you will end up in the right place to catch the ball.

Notice the difference between the 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 tells the body where to go. There is a nice 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, on-going, real-time adjustments to the run, a kind of *co-ordination* between the inner and the outer. Such co-ordination dynamics constitute something of a challenge to traditional ideas about perception and action. For they replace the notion of rich internal representations and computations, with the notion of less expensive strategies whose task is not to first represent the world and then reason on the basis of the representation, so much as to maintain a kind of adaptively potent equilibrium which couples the agent and the world. Whether such strategies are genuinely non-representational and non-computational, or suggestive instead of different *kinds* of representation ("action-oriented representations") and more efficient forms of computation, is a difficult question whose resolution remains uncertian[2,6,7,18,33,34,35].

What is clear, however, is that tuning to higher order invariants can help explain a wide variety of adaptive responses, including visually guided locomotion[36,37], rhythmic movement[18,38], and the capacity to grasp and wield objects[7] (hammer, golf clubs, and so on). In all these cases (see also Box 1) behavioral success involves locking on to simple (but often far from obvious) properties of the environment made available in the perceptual array.

Beyond Adaptive Coupling?So just what *does* the embodiment boom mean for cognitive science? The honest truth is, it is just too soon to tell. But 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 story contribute to our understanding of so-called call "representation-hungry"[39] problem-solving? To get a fix on this idea, consider the much simpler notion of *adaptive coupling*. Adaptive coupling occurs when some system (typically a plant or animal) evolves a mechanism which allows it to track the behavior of some other system (a predator, perhaps, or a source of food or energy). To borrow an example from Brian Cantwell Smith[40], the sunflower has evolved to track the daily motion of the sun across the sky. 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 harbor internal representations? A common – and I think correct – intuition is that it does not. There is nothing *cognitive* in play. One reason we think this is that cognition has long been taken to involve the capacity to relate to an *intentional object*[41] – and this means, in part, an object that may not be present-at-hand and that may not even exist. The sunflower, by contrast, tracks the sun only when the sun is (in a certain sense, at least) actually there. Or more precisely, when there is an ongoing external physical trace to which the sunflower can adaptively couple. The mark of the cognitive, then, is rather plausibly the capacity to engage in something like off-line reason[42] – reasoning in the absence of that which our thoughts concern. Classical ('disembodied') cognitive science had a nice story to tell here, for by positing an inner realm richly populated with internal tokens that *stood for* external objects and states of affairs, 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 Gibsonia approach here, see Kirsh[43] and Van Leeuwen[44]).

One promising move, at this point, is to suggest 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 environmental de-coupleability. The best developed version of this strategy is probably the mobile robotics work of Lynn Andrea Stein at the MIT Artificial Intelligence lab. Stein[45] uses as her platform a mobile robot (designed and implemented by Maja Mataric[46]) named TOTO (see box 2).

TOTO was adept at interacting with the local environment, and could even, in a weak sense, 'track' that which was not present-to-hand: it could return, on command, to a previously encountered location. TOTO could not, however, be prompted to track or 'think about' any location that it had not previously visited. METATOTO[45] 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 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. 45, p. 404).

METATOTO uses the basic behavior-producing architecture of TOTO (see box 2), 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 were the robot

locomoting along a given route on the map. The map can thus induce sequences of experiences which are qualitatively similar to those generated by real sensing and acting, and this allows METATOTO to profit from "virtual experiences', just as TOTO profits (see box 2) from real experience. Once the sensors and motors are restored to real world input and action, 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 story 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 kind of mild dilemma. For the simulation-based story looks most clearly different (from traditional stories involving inner world models) just 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. 45, p. 396). To cash out this claim, Stein reminds us that METATOTO works by simulating both sensors and actuators, and that simulation runs recreate the kinds of feedback (short and long sonar signals, etc.) that would be received from the actual world, were the robot to really change position. There are, of course, some idealizations: the simulated motion is, for example, straight and precise, unaffected by the dinks and slopes found in the real world. But, overall, METATOTO does indeed rely on the simulation of sensori-motor 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 U.S. gun manufacturers should be held legally liable for having knowingly manufactured more guns than the legal market could possibly account for? Here it is not

clear how rich sensori-motor 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 introspectively seem. Rather, it is hard to see how sensori-motor simulation could in principle account for all the kinds of thought and reasoning that the topic demands. Simulated acting and sensing may well play a role-and perhaps even an essential role[4]-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 fundamental moral debate (using concepts like 'liability', reasonable expectation', 'acceptable risk', etc., etc.) does not lend itself-on the face of things- to an analysis in terms of the endogenous generation of, and response to, anything like sequences of low-level sensory inputs and/or motor outputs of the kind involved in *literal* sensori-motor simulation as done by METATOTO. No doubt there are other, less direct ways to depict high level cognition as something like a de-coupled version of sensori-motor coordination strategies, or at least as involving the re-activation, off-line, of aspects of sensory experience[47]. But (and here's the mild dilemma I promised) it does seem that the more decoupled and abstract the target contents become, *either* the less applicable the sensori-motor simulation strategy is, *or* the less clearly it can then be differentiated from the more traditional approaches it seeks to displace.

Simple vs. 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 worry, I want to distinguish two different ways to appeal to facts about embodiment and environmental embedding. The first, which I'll call *"simple embodiment"*, treats such facts as, primarily, *constraints upon a theory of inner organization and processing*. The second, which I'll 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 trades heavily in 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 more. Attention to the roles of body, world and action, in such cases, is basically a methodological tool aimed at getting the *internal* data-structures and operations *right*. Churchland et al's (ref 23, p. 60) vision of a "motocentric" rather than "visuocentric" cognitive science has, I suspect, a similar goal. Maja Mataric's[46] work on TOTO (box 2) has this flavor, insofar as it concentrates attention on an inner representational resource (the map/controller) and is exploring the ways usefulness in real-world action guidance may constrain and inform the nature of inner representations and processing. Ditto, finally, for the recent body of important work on the role of bodily *metaphors*[4] 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, lies in the striking claims involving 'radical embodiment'. Such claims can be found in work by Tim Van Gelder[48], Thelen and Smith[6], Kelso[18], Varela et al[1], Turvey and Carello[7] and others. These visions 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[6,18,48] (see box 3)

II. that traditional notions such as internal representation and computation are inadequate and unnecessary[6,7,48]

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 which cross-cut 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 just) 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[6,30.31] (and see box 1), adult motor actions[7,18], and mobile robotics[9,10,11,21]. The support is weak, however, because the solutions that look most non-computational, representation-free, and dynamical-analysis (see box 3) friendly usually involve cases of (mere) adaptive coupling, and do not directly confront 'representation-hungry' problems. Here, we must simply suspend judgement and await empirical advance.

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 comes into focus. For these new kinds of internal representation may differ from familiar forms both in their *contents* (being more "deictic"[25] and action-oriented[2]) and in the nature of their inner *vehicles* (perhaps using temporally extended processes and complex dynamical regularities as inner 'tokens'[2,49,50]).

A live prospect for the fans of radical embodiment may, however, lie in claim III (the point about 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's[25] use of a notion of "deictic pointers". A pointer, in A. I., 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 all describe them, are physical actions – such as foveating a certain location – which play a similar kind of functional role. The very act of foveation, it is suggested, may be used to temporarily bind color to location, or to direct a reaching motion to a target. In such cases, the author's comment, "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 may likewise prove fruitful to consider the literal extension of the cognitive system to include aspects of the local environment. In this vein, Clark [2,51] and Hutchins [15], following Vygotsky [52], Bruner [53], Dennett [54] 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'[34,55,56,57] human performance. 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 command[57]. 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[52,54]. The external environment, actively structured by us, becomes a source of cognition-enhancing 'wideware'[51]: external items (devices, media, notations) that scaffold and *complement* (but typically do not replicate) biological modes of computation and processing, creating extended cognitive systems whose computational profiles are quite different from those of the naked brain. Hutchins for example,[15] gives a wonderful 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 navigation. 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 institutions[58]) may be for us what the actively created whorls and vortices are for the tuna. Human brains, raised in this sea of cultural tools[54] will 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 exploits the presence of external symbols and problem solving resources, and unraveling the ways biological brains couple themselves with these very special kinds of ecological objects, is surely one of the most exciting tasks confronting the sciences of embodied cognition-and

one that may shed great light on the role of embodiment in more abstract cognitive domains.

Conclusions

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

But many questions remain unanswered. The embodied approach itself seems to come in two distinct varieties. The first ("simple embodiment") stresses the role of body, world and action in informing and constraining stories that still focus on inner computation, representation and problem-solving. The second ("radical embodiment") sees such a *profound* interplay of brain, body and world as to fundamentally transform both the subject matter and the theoretical framework of cognitive science itself.

The major challenge for the vision of 'radical embodiment' 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 thus impact only our ideas about low-level sensorimotor processes. The real lesson seems to be that (in the human case, at least) we find, *at all levels*, a curious mixture of highly "embodied, embedded" 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 resource[40]. (Human language skills – a topic I have deliberately surpressed

in this review – are a case in point, words and text being both real, external objects that we can encounter and manipulate *and* key instruments of inner, abstract, environmentally-decoupled reason[54,59]).

In sum, the gulf between the embodied, embedded skills of the 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 here-and-now 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. The path from fish to fantasy remains faint, but it is still the most promising trail in town.[*]

[*] Thanks to Tim Van Gelder, Esther Thelen, Dan Dennett, Randy Beer, Michael Wheeler, Lynn Stein, Brian Cantwell Smith, Dominic Murphy, Timo Jarvilehto and two anonymous referees for invaluable help and suggestions on these topics. Any errors, misrepresentations or undue conservatism are all my own.

Box 1. The Situated Infant.

Thelen and Smith[a] show in some detail that infant walking abilities depend on complex interactions between neural states, the biomechanics of the legs, and local environmental parameters. 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, p. 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 (ref. a, p. 111-112) seems to be that the stepping behavior depends heavily on a spring-like biomechanical response. To uncoil the spring and shoot 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).

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

Box 2. Maja Mataric's TOTO[a]

TOTO (see fig) uses ultrasonic range sensors to detect walls, corridors, and so on and is able to use its physical explorations to build an inner map of its environment, which it can 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* (ref. 2, p. 49) way, combining information about the robot's movement and correlated perceptual input. The 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 turns out to be a single computational task.

Fig TOTO here

Box 3. Dynamical Systems Theory

The dynamical systems approach has been gaining ground in cognitive scientific treatments of cognition and adaptive behavior[a,b,c,d]. The dynamical approach focuses on the evolution of a system over time, and is particularly well-suited to dealing with cases in which some system or component **a** is constantly affecting and being affected by some system or component **b** (which may likewise be continuously sensitive to some item **c** and so on). An example might be, for example, the process of returning a tennis serve: the locations of the ball and the other player (and perhaps your partner, in doubles) are constantly changing, and all the while, you are moving and acting, which is affecting the other players, and so on; in sum, "everything is simultaneously affecting everything else" (ref. a, p. 23). Such densely coupled unfoldings, 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[e] is expressly designed to deal with such simultaneous interactive complexity. It provides a set of mathematical and conceptual tools that support a basically *geometric* understanding of the space of possible total system behaviors. Such analyses have proven useful in understanding the activity of simple robots[f], infants[d], and adults[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[b] and pursued by Melanie Mitchell and her colleagues[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[h].

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", 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?
- If we follow the embodied, embedded approach to its natural conclusions, how do we avoid losing sight of the distinction between the agent and the world in which she thinks and acts? Where does the mind stop and the rest of the world begin?

References

1 Varela, F., Thompson E., and Rosch, E. (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.), 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 R. Port and T. Van Gelder (eds) *Mind as Motion* (MIT Press) 373-401

8 Damasio, A. (1994) *Descartes' Error*, Grosset

9 Chiel, H. and Beer R. (1997) The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment *Trends in Neurosciences* 20, 553-557

10 Arkin, R. (1998) *Behavior-Based Robotics*, MIT Press

11 Brooks, R. (1991) Intelligence without representation *Artificial Intelligence* 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* 25:4:5-11

14 Greeno, J (1997) "On Claims That Answer the Wrong Questions" (Response to Anderson, et al, ref. 13) Educational Researcher 25:4:5-11

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 *Scientific American* 272, 64-71

21 Raibert and Hodgins (1993) Legged Robots in *Biological Neural Networks in Invertabrate Neuroethology and Robotics* (Beer R. D., Ritzmann R. E. and McKenna, T., eds), pp 319-54, Academic Press

22 Ballard, D. (1991) Animate Vision *Artificial Intelligence* 48, 57-86

23 Churchland, P., Ramachandran V., 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., Hayhoe, M., Pook, P. and Rao, R. (1997) "Dieictic Codes For the Embodiment of Cognition" *Behavioral and Brain Science* 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 W. Collins (ed) *Minnesota Symposium on Child Psychology 15* (Huirdale) p. 55-81

29 Michaels, C. and Cavello, C. (1981) *Direct Perception* (Prentice Hall)

30 Adolph, K (1995) "Psychophysical Assessment of Toddlers' Ability to Cope with Slopes*" Journal of Experimental Psychology: Human Perception and Performance* 21:4:734-750

31 Lee, D. and Reddish, P. (1981) "Plummeting Gannets: A Paradism of Ecological Optics" *Nature* 293:293-294

32 McBeath, M., Shaffer, D. and Kaiser, M. (1995) "How Baseball Outfielders Determine Where to Run to Catch Fly Balls" *Science*:268:569-573

33 Clark A. (1995) “Moving Minds: Situating Content in the Service of Real-time Success” in J. Tomberlin (ed) *Philosophical Perspectives9:A-I, Connectionism and Philosophical Psychology* (Ridgeview) 89-104

34 Rutkowska, J. (1993) *The Computational Infant* (Harvester)

35 Neisser, U (1976) *Cognition and Reality* (Freeman)

36 Adolph, K and Eppler, M. (1998) “Development of Visually Guided Locomotion” *Ecological Psychology* 10(3-4) 303-321

37 Adolph, K, Eppler, M. and Gibson, E. (1993) “Crawling versus walking Infants’ perception of affordance for locomotion over sloping surfaces” *Child Development* 64:1158-1174

38 Hatsopoulis, N. and Warren, W. (1996) “Resonance Tuning in Rhythmic Arm Movements” *Journal of Motor Behavio*r 78:1:3-14

39 Clark, A. and Toribio, J. (1995) Doing Without Representing? *Synthese* 101: 401-31

40 Smith, B. C. (1996) *On the Origin of Objects,* MIT Press

41 Brentano, F. (1874) *Psychologie vom Empirischen Standpunkt*, translated as *Psychology from an Empirical Standpoint* (1973) London

42 Clark, A. and Grush, R. Towards a Cognitive Robotics *Adaptive Behavior* (in press)

43 Kirsh, D. (1992) “Today the Earwig, Tomorrow Man?” in D. Kirsh (ed) Foundations of Artificial Intelligence (MIT Press) 161-184

44 Van Leeuwen, C (1998) “Perception” Entry in W. Bechtel and G. Graham (eds) *A Companion to Cognitive Science* (Blackwell) 265-281

45 Stein, L. (1994) Imagination and Situated Cognition *Journal of Exp Theor Artif Intell* 6, 393-407

46 Mataric, M. (1992) Integration of representation into goal-driven behavior based robots *IEEE Journal of Robotics and Automation* 8, 304-12

47 Barsalou, L. Perceptual symbol systems *Behavioral and Brain Science* (in press)

48 van Gelder, T. (1995) What Might Cognition Be, If Not Computation? *Journal of Philosophy* XCII, 345-381

49 Clark, A. (1997) The Dynamical Challenge *Cognitive Science* 21, 461-481

50 Mitchell, M., Crutchfield, J. and Hraber, P. (1994) Evolving Cellular Automata to Perform Computations *Physica D* 75, 361-91.

51 Clark, A. (1998) Where Brain, Body and World Collide *Daedalus* 127, 257-80

52 Vygotsky, L. (1986) *Thought and Language,* MIT Press

53 Bruner, T (1990) *Acts of Meaning* (Harvard)

54 Dennett, D. (1996) *Kinds of Minds* (Basic Books)

55 Wood, D., Bruner, J. and Russ G. (1976) "The Role of Tutoring in Problem Solving" Journal of Child Psychology and Psychiatry 17, 89-100

56 Norman, D. A. (1993) Things That Make Us Smart (Addison Wesley)

57 McClelland, J., D. Rumelhart, et al. (1986) The Appeal of Parallel Distributed Processing, in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (McClelland, J. , Rumelhart, D. and PDP Research Group, eds), vol II, pp 3-44, MIT Press/Bradford Books

58 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

59 Clark, A. (1998) Magic Words: How Language Augments Human Computation, in *Language and Thought* (Boucher, J. and Carruthers, P., eds), pp. 162-83, Cambridge University Press

References (Box 1)

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 *Cognitive Science* 21, 461-81

References (Box 2)

a Mataric, M.J. (1992) Integration of representation into goal-driven, behavior-based robots *IEEE Journal of Robotics and Automation* 8, 304-12

References (Box 3)

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 *Artificial Intelligence* 72, 173-215

g Mitchell, M., Crutchfield, J., and Hraber, 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 *Proceedings of the National Academy of Sciences* 92, 10742-10746