7 DYNAMICS

7.1 Sketches

Cognitive science, we have seen, is involved in an escalating retreat from the inner symbol: a kind of inner symbol flight. The original computational vision (Chapters 1 and 2) displayed no such qualms and happily tied syntax to semantics using static inner items that could stand for semantic contents. Such items were invariant ("token identical") across different contexts and were easily thought of as inner symbols. Connectionist approaches (Chapter 4) expanded our conception of the syntax/semantics link, allowing context-sensitive coalitions of unit activity to bear the semantic burden and producing sensible behaviors and judgments without the use of static, chunky, easy-to-interpret inner states. Connectionism, we might say, showed us how to believe in internal *representations* without quite believing in traditional internal *symbols*. Recent work in neuroscience, animate vision, robotics, and artificial life (Chapters 5 and 6) has expanded our conceptions still further, by displaying an even wider range of neural dynamics and possible coding strategies and by stressing the profound roles of timing, body, motion, and local environment in biological problem solving.

But as the complexity and environmental interactivity of our stories increase, so the explanatory leverage provided by the original complex of theoretical notions (symbols, internal representations, computations) seems to diminish. Dynamic systems theory, as it is used in recent[1] cognitive science, can be seen as an attempt

[1]Dynamic approaches to cognition go back at least as far as the wonderful cybernetics literature of the 1940s and 1950s—see, e.g., Wiener (1948) and Ashby (1952, 1956). But the approach fell into disfavor in the early days of symbol system A.I. Its recent resurgence owes a lot to the efforts of theorists such as Kelso (1995), van Gelder (1995),

to find analytic tools better suited to the study of complex interactive systems. Whether such tools offer an out-and-out *alternative* to the traditional theoretical framework, or are better seen as a kind of subtle *complement* to that framework, are matters to which we will soon return. The first order of business is to clarify what a dynamic approach involves.

Dynamic systems theory is a well-established framework in physical science.[2] It is primarily geared to modeling and describing phenomena that involve change over time (and change in rate of change over time, and so on). Indeed, the broadest definition of a dynamic system is simply any system that changes over time. Just about every system in the physical world (including all computational systems) is thus a dynamic system. But it is only when the patterns of change over time exhibit a certain kind of complexity that the technical apparatus of dynamic systems theory really comes into its own. Some of the key features on which this special kind of explanatory power depends include

1. the discovery of powerful but low-dimensional descriptions of systemic unfolding,
2. the provision of intuitive, geometric images of the state space of the system,
3. the (closely related) practice of isolating *control parameters* and defining *collective variables* (see below), and
4. the use of the technical notion of *coupling* (see below) to model and track processes involving continuous circular causal influence among multiple subsystems.

Transposed into the cognitive scientific domain, these features make dynamic approaches especially attractive for understanding those aspects of adaptive behavior that depend on complex, circular causal exchanges in which some inner factor x is continuously affecting and being affected by some other (inner or outer) factor y (which may itself stand in similar relations to a factor z, and so on). Such complex causal webs, as we began to see in the previous chapter, are often characteristic of natural systems in which neural processing, bodily action, and environmental forces are constantly and complexly combined. To get the flavor of the dynamic approach in action, let us review a few examples.

Case 1: Rhythmic Finger Motion

Consider the case (Kelso, 1981, 1995, Chapter 2) of rhythmic finger motion. Human subjects, asked to move their two index fingers at the same frequency in a side-to-side "wiggling" motion, display two stable strategies. Either the fingers move in phase (the equivalent muscles of each hand contract at the same moment), or exactly antiphase (one contracts as the other expands). The antiphase solution, however, is unstable at high frequencies of oscillation—at a critical frequency it collapses into the phased solution.

How should we explain and understand this patter of results? One strategy is to seek a more illuminating description of the behavioral events. To this end, Kelso and his colleagues plotted the phase relationship between the two fingers. This variable is constant for a wide range of oscillation frequencies but is subject to a dramatic shift at a critical value—the moment of the antiphase/phase shift. Plotting the unfolding of the relative phase variable is plotting the values of what is known as a "collective variable," whose value is set by a *relation* between the values of other variables (the ones describing individual finger motions). The values of these collective variables are fixed by the frequency of motion, which thus acts as a so-called control parameter. The dynamic analysis is then fleshed out by the provision of a detailed mathematical description—a set of equations displaying the space of possible temporal evolutions of relative phase as governed by the control parameter. This description fixes, in detail, the so-called state space of the system. A systemic state is defined by assigning a value to each systemic variable, and the overall state *space* (also known as phase space) is just the set of all possible values for these variables—all the value combinations that could actually come about. Dynamicists often think about target systems in terms of possible trajectories through such state spaces—possible sequences of states that could take the system from one location in state space to another. The set of possible trajectories through a state space is called the "flow." Finally, certain regions of the state space may exhibit notable properties (see Box 7.1). An *attractor* is a point or region such that any trajectory passing close by will be drawn into the region (the area of such influence being known as the basin of attraction). A *repellor* is a point or region that deflects incoming trajectories. A *bifurcation* occurs when a small change in parameter values can reshape the flow within the state space and yield a new landscape of attractors, repellors, and so on. Dynamic systems approaches thus provide a set of mathematical and conceptual tools that helps display the way a system changes over time.

In the case of rhythmic finger motion, Haken, Kelso, and Bunz (1985) use a dynamic analysis to display how different patterns of finger coordination (in-phase/antiphase, etc.) result from different values of the control parameter (frequency of oscillation). This detailed dynamic model was capable of (1) accounting for the observed phase transitions without positing any special "switching mechanism"—instead, the switching emerges as a natural product of the normal, self-organizing evolution of the system, (2) predicting and explaining the results of selective interference with the system (as when one finger is temporarily forced out of its stable phase relation), and (3) generating accurate predictions of, e.g., the time taken to switch from antiphase to phase. For a nice review of the model, see (Kelso, 1995, pp. 54–61).

A good dynamic explanation is thus perched midway between what, to a more traditional cognitive scientist, may at first look like a ("mere") *description* of a pattern of events and a real *explanation* of why the events unfold as they do. It is not a mere description since the parameters need to be very carefully chosen so that the resulting model has predictive force: it tells us enough about the system to know

Box 7.1

Numerical Dynamics

Stan Franklin (1995), in his *Artificial Minds* (Chapter 12), offers a useful introductory example of a dynamical analysis. Consider the real numbers (with infinity). And imagine that the global dynamics of the number space are set by a squaring function so that for any number x, given as input (initial state), the next state of the system will be x^2. Now consider what happens assuming different initial states. If the initial state is the input 0, the numerical unfolding stays at 0: this is an example of "converging to a fixed point attractor." For initial state 2, the numerical unfolding continues 4, 16, 256, converging to infinity. For initial state −1, the system goes to 1 and stops. But initial points close to 1 (0.9, etc.) move rapidly away. 0 and infinity are thus the attractors to which many initial states converge. 1 is a repellor; a point from which most initial states move away. Since all initial states between 0 and 1 head progressively toward 0 (as the numbers become smaller and smaller with each application of the squaring function), 0 has a basin of attraction that includes all these points (in fact, all the real numbers between −1 and 1). Infinity has a basin of attraction that includes all points greater than 1 or less than −1. The general situation is illustrated in Figure 7.1. To produce so-called periodic behavior, it is necessary to alter the global dynamics, e.g., to the square of the input number, minus 1. An initial state of −1 will then display the repeating (periodic) trajectory: 0, −1, 0, −1, etc.

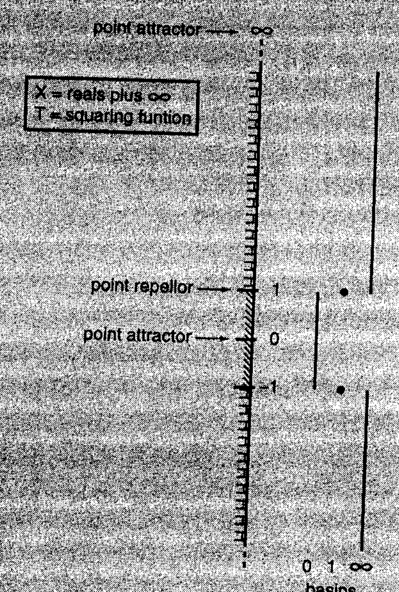

Figure 7.1 Basins of attraction. [From Franklin (1995, p. 283). By permission of MIT Press.]

how it would behave in various nonactual circumstances. But it differs from more traditional cognitive scientific explanations in that it greatly abstracts away from the behavior of individual systemic components.

Case 2: Treadmill Stepping[3]

Consider the phenomenon of learning to walk. Learning to walk involves a regular pattern of developmental events that includes (1) the ability, present at birth, to produce coordinated stepping motions when held upright in the air, (2) the disappearance, at about 2 months, of this response, (3) its reappearance at around 8–10 months when the child begins to support its own weight on its feet, and (4) the appearance, at around 1 year old, of independent coordinated stepping (walking).

At one time, it was thought that the best explanation of this overall pattern would depict at as the expression of a prior set of instructions, complete with timing, encoded in (perhaps) a genetically specified central pattern generator (see Thelen and Smith, 1994, pp. 8–20, 263–266). Thelen and Smith (1994) argue, however, that there is no such privileged, complete, and prespecified neutral control system, and that learning to walk involves a complex set of interactions between neural states, the spring-like properties of leg muscles, and the local environment. Walking, according to Thelen and Smith, emerges from the balanced interplay of multiple factors spanning brain, body, and world, and is best understood using a dynamic approach that charts the interactions between factors and that identifies crucial elements on "control parameters."

Thelen and Smith conducted a fascinating sequence of experiments yielding broad support for such a view. Two especially significant findings were

1. that stepping motions can be induced during the "nonstepping" window (2–8 months) by simply holding the baby upright in warm water (instead of air) and
2. that nonstepping 7 month olds held upright on a motorized treadmill perform coordinated alternating stepping motion, and are even able to compensate for twin belts driving each leg at a different speed!

The explanation, according to Thelen and Smith (1994, Chapters 1 and 4), is that stepping is dynamically assembled rather than being the expression of a simple inner command system. Bodily parameters such as the leg weight, which is effectively manipulated by partial immersion in water, and environmental factors (such as the presence of the treadmill) seem equally implicated in the observed behaviors. In the case of the treadmill, further experiments revealed that the crucial

[3]This case is treated in more detail in Clark (1997).

factor was the orientation of leg and foot to the treadmill. Infants who made flat-foot belt contact exhibited treadmill stepping, whereas those that made only toe contact failed to step. Thelen and Smith (1994, pp. 111–112) explain this by hypothesizing that the infant leg, when stretched out, is acting like a spring. At full back stretch, the spring uncoils and swings the leg forward. Flat-foot belt contact precociously ensures this full back stretch and hence initiates stepping. Relative flexor or extensor tendencies in the legs thus contribute heavily to the emergence of coordinated stepping in the normal case (Thelen and Smith, 1994, p. 113). The treadmill stepping task provides an especially useful window onto the dynamic construction of infant walking, as it highlights the complex and subtle interplay between intrinsic dynamics, organic change, and external task environment. In dynamic terms, the treadmill looks to be acting as a real-time control parameter that prompts the phase shift, in 7 month olds, from nonstepping to smooth alternating stepping motions. Stepping behavior thus "emerges only when the central elements cooperate with the effecters—the muscles, joints, tendons—in the appropriate physical context" (Thelen and Smith, 1994, p. 113).

Case 3: The Watt Governor

Consider finally a classic example recently deployed by Tim van Gelder (1995)—the operation of the Watt (or centrifugal) governor. The job of the governor is to keep constant the speed of a flywheel that drives industrial machinery and is itself driven by a steam engine. Given variations in steam pressure and current workload (number of machines being driven, etc.), the flywheel speed tends to fluctuate. To keep it smooth and constant, the amount of steam entering the pistons is controlled by a throttle valve. More steam results in more speed; less steam results in less speed. At one time, a human engineer had the unenviable task of making these constant corrections. How might such a process be automated?

One solution (which van Gelder describes as the computational solution) would involve a sequence of steps and measurements. For example, we might program a device to measure the speed of the flywheel, compare this to some desired speed, measure the steam pressure, calculate any change in pressure needed to maintain the desired speed, adjust the throttle valve accordingly, then begin the whole sequence anew (see van Gelder, 1995, p. 348). What makes this kind of solution *computational,* van Gelder suggests, is a complex of familiar features. The most important one is representation: the device measures the speed of the flywheel, creates a token that stands for the speed, and performs numerous operations (comparisons, etc.) on this and other representations. These operations are discrete and occur in a set sequence, which then repeats itself. The sequence involves a perception/measurement–computation–action cycle in which the environment is measured ("perceived"), internal representations created, computations performed, and an action chosen. The overall device reflects a nicely decomposable problem solution. For it respects a division of the problem into these distinct subparts, each of which is dealt with independently, and which are coordinated by acts of communication (in which *x* tells *y* the value of *z* and so on). The features distinctive of the computational governor are thus (1) the use of internal representations and symbols, (2) the use of computational operations that alter and transform those representations, (3) the presence of a well-defined perception–computation–action cycle (what van Gelder calls "sequential and cyclic operation"), and (4) the susceptibility to step-wise information-processing decomposition (what van Gelder calls "homuncularity").

Now for the second solution, the one discovered by James Watt (see Figure 7.2). Gear a vertical spindle into the flywheel and attach two hinged arms to the spindle. To the end of each arm, attach a metal ball. Link the arms to the throttle valve so that the higher the arms swing out, the less steam is allowed through. As the spindle turns, centrifugal force causes the arms to fly out. The faster it turns, the higher the arms fly out. But this now reduces steam flow, causing the engine to slow down and the arms to fall. This, of course, opens the valve and allows more steam to flow. By clever calibration this centrifugal governor can be set up so as to maintain engine speed smoothly despite wide variations in pressure, workload, and so on. (This story is condensed from van Gelder, 1995, pp. 347–350.)

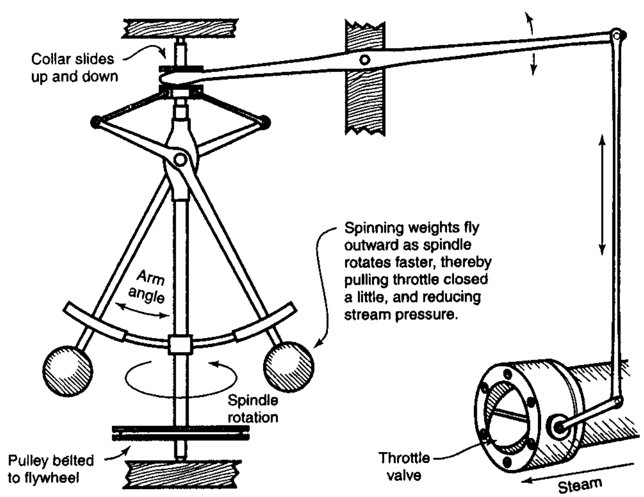

Figure 7.2 The Watt centrifugal governor for controlling the speed of a steam engine (Farey, 1827). [From van Gelder, T.J. (1997). "Dynamics and cognition." In J. Haugeland, ed., *Mind Design II: Philosophy, Psychology, and Artificial Intelligence*, rev. ed. Cambridge, MA: MIT Press. Reproduced by kind permission of the author and the publishers, MIT Press.]

This centrifugal governor, van Gelder claims, constitutes a control system that is noncomputational, nonrepresentational, and that simply cries out for dynamic analysis and understanding. In particular, only a dynamic analysis can explain the peculiarly complex, yet effective, relationship that is obtained between the arm angle and the engine speed. A mad-dog representationalist might, perhaps, try to claim that the best way to understand this relationship is by depicting the arm angle as a representation of the engine speed. But, van Gelder (1995, p. 353) insists, the real relationship is "much more subtle and complex than the standard notion of representation can handle." It is more subtle and complex because the arm angle is continuously modulating the engine speed at the same time as the engine speed is modulating the arm angle. The two quantities are best seen as being codetermined and codetermining—a relationship nicely captured using a dynamic apparatus (see below) of coupled differential equations. The Watt governor then fails to constitute a *computational* device for two reasons. First, because, on van Gelder's (1995, p. 353) account, computation requires the manipulation of token-like representations that seem notably absent here. And second, because there are no discrete operations in the governing processes and hence no distinct sequence of manipulations to identify with the steps in an algorithmic solution. The Watt governor thus fails to exhibit any of the features associated with the computational solution, and for a single deep reason: the continuous and simultaneous relations of causal influence that obtain among the various factors involved. It is this distinctive kind of causal profile that both invites treatment in terms of an alternative dynamic analysis and that causes problems for the traditional (computational and representational) approach.

The way to capture such a complex causal relationship, van Gelder asserts, is by using the dynamic notion of *coupling*. In a typical quantitative dynamic explanation, the theorist specifies a set of parameters whose collective evolution is governed by a set of differential equations. Such explanations allow distinct components (such as the arm and the engine) to be treated as a coupled system in a specific technical sense, viz. the equation describing the evolution of each component contains a slot that factors in the other one's current state (technically, the state variables of the first system are the parameters of the second and vice versa). Thus consider two wall-mounted pendulums placed in close proximity on a single wall. The two pendulums will tend (courtesy of vibrations running along the wall) to become swing-synchronized over time. This process admits of an elegant dynamic explanation in which the two pendulums are analyzed as a single coupled system with the motion equation for each one including a term representing the influence of the other's current state.[4] This kind of complex, constant, mutual interaction is, van Gelder and others claim,[5] much closer to the true profile of agent–

[4]See Salzman and Newsome (1994).

[5]For example, Beer and Gallagher (1992) and Wheeler (1994).

environment interactions than is the traditional vision of a simple perception–computation–action sequence.

With these comments and case studies in hand, it is now reasonably easy to construct the case for a dynamic cognitive science. The case turns on three basic assertions.

The first, relatively unproblematic, assertion is that body and world (and hence time, movement, and so on) all matter, and can play powerful roles in adaptive problem solving. We have seen several examples of this earlier in the text, e.g., the work on infant locomotion, cricket phonotaxis, and animate vision, as well as in a wealth of research in biology, cognitive neuroethology, and robotics.[6] The second assertion is that body and world matter not simply because they provide an arena for useful action and a sensitive perceptual front-end, but because neural, bodily, and environmental elements are intimately intermingled courtesy of processes of continuous reciprocal causation that criss-cross intuitive boundaries. This leads to the third and final assertion, that the traditional tools of computational and representational analysis (with the associated image of an input–compute–act cycle) cannot do justice to such a complex interactive process and that the mathematical and topological resources of dynamic systems theory are to be preferred. Such, it seems to me, is the central argument.[7] But is it really powerful enough to win the day?

7.2 Discussion

A. THE HIDDEN PREMISE

The most radical conclusion to be drawn from the dynamic considerations seems to go something like this:

> *The Radical Embodied Cognition Thesis*
> Structured, symbolic, representational, and computational views of cognition are mistaken. Embodied cognition is best studied using noncomputational and nonrepresentational ideas and explanatory schemes, and especially the tools of dynamic systems theory.

[6]For review, see Clark (1997).

[7]The centrality of the point about continuous reciprocal causation is evident from remarks such as these: "the . . . deepest reason for supposing that the centrifugal governor is not representational is that . . . arm angle and engine speed are at all times both determined by, and determining each other's behavior. [This relationship] is much more subtle and complex than the standard concept of representation can handle" (van Gelder, 1995, p. 353). Or again: "adaptive behavior is the result of the continuous interaction between the nervous system, the body and the environment . . . one cannot assign credit for adaptive behavior to any one piece of this coupled system" (Chiel and Beer, 1997, p. 555). See also van Gelder and Port (1995, pp. ix, 23), Schöner (1993), Kelso (1995), and the discussion in Clark (1998c).

Given the nature of the dynamic demonstrations, it seems initially surprising to find such radical and sweeping conclusions. What we seem to have before us is, surely, just an argument that some quite low-level sensorimotor engagements with the world (finger wiggling, infant walking, Watt governing, etc.) exhibit a complex causal structure that makes it hard to fully explain such engagements using standard notions of computations and representation, and the input–compute–act cycle. This seems compatible with (1) the idea that for *higher level* cognition, the standard framework is still the best and (2) the idea that even at the lower levels, *some aspects* of systemic unfolding might still reward a more traditional analysis.

Despite this, there can be little doubt that genuine and sweeping radical reform is in the air. Thelen and Smith clearly support the radical thesis, writing that:

> *Explanations in terms of structure in the head*—beliefs, rules, concepts and schemata—*are not acceptable*. . . . Our theory has new concepts at the center—nonlinearity, reentrance, coupling heterochronicity, attractors, momentum, state spaces, intrinsic dynamics, forces. These concepts are not reducible to the old ones. (Thelen and Smith, 1994, p. 339; my emphasis)

> We posit that development happens because of the time-locked pattern of activity across heterogenous components. We are not building representations of the world by connecting temporally contingent ideas. *We are not building representations at all! Mind is activity in time* . . . the real time of real physical causes. (Thelen and Smith, 1994, p. 338; my emphasis)

Scott Kelso, though more sympathetic to a (reconceived) notion of internal information bearers (representations?), asserts that

> The thesis here is that the human brain is *fundamentally* a pattern-forming, self-organized system governed by non-linear dynamical laws. *Rather than compute*, our brain dwells (at least for short times) in metastable states. (Kelso, 1995, p. 26; second emphasis mine)

Other writers who sometimes seem tempted by the radical thesis include Smithers (1994), Wheeler (1994), Maturana and Varela (1980), Skarda and Freeman (1987), and van Gelder (1995). The generally balanced and extended treatment in Keijzer (1998, p. 240) also leans toward the radical conclusion, suggesting that attempts (such as Clark, 1997) to preserve the use of a computational/representational framework amount to "the injection of a particular set of thought habits into a tentative and still fragile interactionist account of behavior."

The first order of business, then, is to somehow join the dots, to identify the additional ideas and premises that might link the rather limited empirical demonstrations to the sweeping radical conclusion. The most crucial linking theme, I now believe,[8] relates to the idea of the continuity of life and mind. We have already encountered this idea (in Chapter 6), so let us be brief.

[8]Thanks to Esther Thelen for insisting (personal communications) on the importance of this idea.

Consider—following Pollack (1994)—the history of flight. When we first encounter birds and wonder how they manage to fly, the most superficially salient feature might seem to be the flapping of wings. But, as we all now know, and as some early pioneers found out by bitter experience, powerful flapping is not really the key. Instead, as the Wright brothers finally figured out:

> most of the problem of flying is in finding a place within the weight/size dimension where gliding is possible, and getting the control systems for dynamical equilibrium right. Flapping is the last piece, the propulsive engine, but in all its furiousness it blocks our perception. (Pollack, 1994, p. 188)

Specifically, what the flapping obscures is the pivotal importance of what is known as the Aileron principle—the use of control cables and raisable and lowerable wing flaps to allow the pilot to balance the machine while gliding in the air.

> The analogical extension to dynamical approaches to cognition is pretty direct: Just like flapping, symbolic thought is the last piece [of the puzzle] . . . in all its furiousness it obscures our perception of cognition as an exquisite control system . . . governing a very complicated real-time physical system. (Pollack, 1994, p. 118)

Understanding that real-time physical system, Pollack believes, is pretty impossible as long as we focus on symbolic problem solving (flapping). Instead, we should (Pollack, 1994, p. 119) "unify cognition with nature"—look not at "software law" but at physical law. Only then will we begin to see how biological intelligence can be as robust and flexible as it is—how, for example, the injured cat can immediately adopt a successful three-legged gait courtesy of the complex, interactive dynamics linking neural nets with spring-like muscle and tendon systems. Such rich interactive dynamics have little, it seems, to do with explicit, symbol-using problem solving. Yet it is this rich nonsymbolic substrate that, it is argued, forms the essential basis for all aspects of biological intelligence (see, e.g., Thelen and Smith, 1994, p. xxiii). This substrate, as we saw, is characterized by processes of continuous reciprocal causal influence in which overall interaction dynamics (rather than some privileged, knowledge-based component) enable the organism to achieve its goals and to compensate for unwelcome environmental changes. It is in this way that the Watt governor, although clearly itself a noncognitive device, may be presented (van Gelder, 1995, p. 358) as "more relevantly similar" in its operation to (the fundamental, dynamical substrate of) human cognition than more traditional computation-and-representation invoking benchmarks such as SOAR (Chapter 2) or even NETtalk (Chapter 4).

B. STRONG AND WEAK CONTINUITY

The radical thesis is rooted, then, in a familiar observation: the shape and operation of higher level cognitive processes have probably been built, in some highly path-dependent fashion, on a more evolutionary basic substrate of perception and sensorimotor control. Connectionists, however (recall Chapter 4) have made similar points, as have theorists working in traditional artificial intelligence (e.g., Simon, 1996), and done so *without* calling into question the fundamental framework of computational and representational explanation. Where's the difference?

The difference again lies in the dynamicist's emphasis on interaction and continuous reciprocal causation; the idea that it is the on-going couplings between environment, body, and nervous system that form the basis of real-time adaptive response. Accepting both path dependence and the interactive nature of basic sensorimotor adaptation, however, *still* falls well short of establishing the thesis of radical embodied cognition.

Thus consider a traditional claim—that we sometimes solve problems by exploiting inner models that are designed (by learning or evolution) to function as off-line *stand-ins* for features of our real-world environment. In such cases, we temporarily abandon the strategy of directly interacting with our world so as to engage in more "vicarious" forms of exploration. It is certainly possible that such off-line problem solving is perfectly continuous with various on-line, highly environmentally interactive, motor control strategies. Thus Grush (1995) describes a piece of circuitry whose principal role is the fine-tuning of on-line reaching. The circuitry, however, involves the use of an inner model (an "emulator loop") that predicts sensory feedback in advance of the actual signals arriving (rather too slowly) from the bodily peripheries. This inner loop, once in place, supports the additional functionality of fully off-line deployment, allowing the system to rehearse motor actions entirely in its "imagination." Such cases suggest both a profound *continuity* between smooth motor control strategies and higher cognitive capacities such as off-line reasoning and imagination, and (simultaneously) a profound *discontinuity* in that the system is now using specific and identifiable inner states as full-blooded stand-ins for specific extraneural (in this case bodily) states of affairs. These are surely internal representations in quite a strong sense. At such times the system is *not* continuously assembling its behavior by balancing ongoing neural bodily and environmental influences. We thus preserve a kind of architectural continuity, but without the added commitment to the radical embodied cognition thesis (for a fuller treatment, see Clark and Grush, 1999).

C. REPRESENTATION AND COMPUTATION, AGAIN

Another worry concerns the nature (content) of any putative internal representations. For it looks as if the target of much dynamicist skepticism is not internal representation per se so much as a particular type of internal representation, viz. what are sometimes called "objectivist" representations—the kind that might be featured in a detailed, viewpoint-independent model of some aspect of the world. Notice, then, a second (and I believe, highly significant—see Clark, 1995) way in which higher level cognition may be continuous with its motor and developmental roots. It may be continuous insofar as it involves internal representations whose contents (unlike detailed "objectivist' representations) are heavily geared toward the support of typical or important kinds of real-world, real-time action. Such contents may (as in the previous example) sometimes be manipulated "off-line"—but they are nonetheless *types* of content (what I elsewhere call action-oriented contents) that are especially suited to the control and coordination of real action in real time. Cognition, on this model, need not always be *actually* interactive (involving brain, body, and world as equal partners). But the inner economy is deeply sculpted and shaped by the interactive needs and patterns of the organism.

Much dynamicist skepticism, on closer examination, looks to address only the notion of objectivist (detached, action-independent, highly-detailed, static, general-purpose) internal representations. Thus Thelen and Smith (1994, pp. 37–44) question all these ideas, suggesting instead that we treat knowledge as an action-guiding process continually organized against a contextual backdrop that brings forth its form, content, and use. The same emphases characterize Varela's notion of "enaction" in which "cognitive structures emerge from the recurrent sensorimotor patterns that enable action to be perceptually guided" (Varela, Thompson, and Rosch, 1991, p. 173). In a related vein, Agre (1995, p. 19) notes the importance of "indexical-functional representations" (such as "a few feet straight ahead")—these are ideal for the cheap control of individual action and are contrasted with objectivist map-like representations such as "at latitude 41, longitude 13." Perhaps, then, some of the disputes really concern the content, not the existence, of inner states whose role is to stand in for important extraneural states of affairs.

Related to this may be an assumption concerning the type of inner control implicated in broadly representationalist/computationalist accounts. The assumption, roughly, is that computational models involve the storage and use of complex inner control structures that plot, in *explicit detail*, all the values and settings of all the physical parameters involved in a given action. Something like this assumption would help explain why Thelen and Smith repeatedly associate the idea that the brain is a computational device with seemingly orthogonal ideas about detailed advance blueprints for behavior, complete with internal clocks, full specifications of all relevant parameter settings (joint-angle coordinates, muscle fixing patterns, etc.) for the limbs, and capable of controlling movement by "'pure' neural commands" (Thelen and Smith, 1994, p. 75, see also pp. xix, 62–63, 71, 264). They then *contrast* this vision of highly detailed, complete neural instruction sets with the ideas of collective states, phase shifts, and control parameters, as discussed earlier. Certain preferred collective states of the system are depicted as synergetic wholes that can be brought forth (but not "programmed") by the action of some control parameter (such as frequency of motion in the rhythmic finger motion case and flexor tone in the treadmill stepping case). Kelso's description of the brain itself as not a computing device but a "pattern-forming, self-organized system" (Kelso, 1995, p. 26) has the same flavor. The contrast is between systems whose behavior is fixed by complete encoded instruction sets and ones whose behavior

emerges as a sequence of temporarily stable states of a complex system with richly interdependent intrinsic dynamics. The slogan may be "patterns without programs"; but the real target is the idea that we use complex neural instruction sets to force orderly behavior from multiple muscles, links, joints, etc. Such detailed forcing is not necessary, it is claimed, because the system self-organizes into a smaller set of preferred states whose flux may be controlled by the action of some simple parameter. (It is a little as if the "computationalist," faced with the problem of moving a crowd from A to B, were to encode an instruction for each person's trajectory, whereas the dynamicist simply finds a control parameter (maybe increasing the heat on one side of the crowd) that then exploits the intrinsic responses of those closest to it, whose motion in turn entrains the movements of their near neighbors, until the crowd moves as a unified whole in the desired direction).

This is an important and fascinating shift in emphasis, to be sure. But does it really amount to a rejection of the idea that the brain computes? I suggest that it cannot, since there is no necessary commitment on the part of the computationalist to the idea of highly detailed or complete instruction sets. A short piece of software, written in a high-level language, will not *itself* specify how or when to achieve many subgoals—these tasks are ceded to built-in features of the operating system or to the activity of a cascade of lower level code. Moreover, a program can perfectly well "assume" a necessary backdrop of environmental or bodily structures and dynamics. Jordan et al. (1994) describe a program for the control of arm motions, but one that assumes (for its success) a lot of extrinsic dynamics such as mass of arm, spring of muscle, and force of gravity.

Now it may be that so very much is done by the synergetic dynamics of the body–environment system that the neural contributions are indeed best treated, at times, as just the application of simple forces to a complex but highly interanimated system whose intrinsic dynamics then carry most of the load. But less radically, it may be that motor activity simply requires *less* in the way of detailed inner instruction sets than we might have supposed, courtesy of the existence of a small set of preferred collective states such that successful behavior often requires only the setting of a few central parameters such as initial stiffness in a spring-like muscle system and so on. Such sparse specifications may support complex global effects without directly specifying joint-angle configurations and the like.

The lack of a particularly detailed kind of neural instruction set does not then establish the total absence of stored programs. Such a characterization is compelling only at the most extreme end of a genuine continuum. Between the two extremes lies the interesting space of what I elsewhere (Clark, 1997) call "partial programs"—minimal instruction sets that maximally exploit the inherent (bodily and environmental) dynamics of the controlled system. The real moral of much actual dynamic systems-oriented research is, I suspect, that it is *in this space that we may expect to encounter nature's own programs.*

D. THE SPACE OF REASONS

The deepest problem with the dynamic alternative lies, however, in its treatment of the brain as *just one more factor* in the complex overall web of causal influences. In one sense this is obviously true. Inner and outer factors do conspire to support many kinds of adaptive success. But in another sense it is either false, or our world view will have to change in some very dramatic fashion indeed. For we do suppose that it is the staggering structural complexity and variability of the brain that are the key to understanding the specifically intelligence-based route to evolutionary success. And we do suppose that that route involves the ability, courtesy of complex neural events, to become appraised of *information* concerning our surroundings, and to use that information as a guide to present and future action. If these are not truisms, they are very close to being so. But as soon as we embrace the notion of the brain as the principal seat of information-processing activity, we are already seeing it as fundamentally different from, say, the flow of a river or the activity of a volcano. And this is a difference that needs to be reflected in our scientific analysis: a difference that typically *is* reflected when we pursue the kind of information-processing model associated with computational approaches, but that looks to be lost if we treat the brain in exactly the same terms as, say, the Watt governor, or the beating of a heart, or the unfolding of a basic chemical reaction.[9]

The question, in short, is how to do justice to the idea that there is a principled *distinction* between knowledge-based and merely physical-causal systems. It does not seem likely that the dynamicist will deny that there is a difference (though hints of such a denial[10] are sometimes to be found). But rather than responding by embracing a different vocabulary for the understanding and analysis of brain events (at least as they pertain to cognition), the dynamicist recasts the issue as the explanation of distinctive kinds of behavioral flexibility and hopes to explain that flexibility using the very same apparatus that works for other physical systems, such as the Watt governor.

Such apparatus, however, may not be intrinsically well suited to explaining the particular way certain neural processes contribute to behavioral flexibility. This is because it is unclear how it can do justice to the fundamental ideas of agency and of information-guided choice. Isn't there a (morally and scientifically) crucial distinction between systems that select actions for reasons and on the basis of acquired knowledge, and other (often highly complex) systems that do not display such goal-oriented behaviors? The image of brain, body, and world as a single, densely coupled system threatens to eliminate the idea of purposive agency unless it is combined with some recognition of the special way goals and knowledge figure in the origination of some of our bodily motions.[11] The computational/information-processing approach provides such recognition by embracing a kind of dual-aspect account in which certain inner states and processes act as the vehicles of knowledge and information.

[9]For the last two cases, see Goodwin (1995, p. 60).

[10]For example, van Gelder's comments (1995, p. 358) on tasks that may only initially appear to require "that the system have knowledge of and reason about, its environment," and Thelen and Smith's (1994, p. xix) stress on the brain as a thermodynamic system. By contrast, the dynamicist Kelso (1995, p. 288) sees the key problem as "how *information* is to be conceived in living things, in general, and the brain in particular."

Perhaps, then, what is needed is a kind of dynamic computationalism in which the details of the flow of information are every bit as important as the larger scale dynamics, and in which some local dynamic features lead a double life as elements in an information-processing economy. Here, then, is one way in which dynamic and computational analyses may proceed hand in hand.[12] The dynamic analyses may help identify the complex and temporally extended physical processes that act as the *vehicles* of representational content. Traditional computationalism may have been just too narrow minded in its vision of the likely syntactic form of the inner bearers of information and content. Our fascination with the static characters and strings of natural language led us to expect simple, local, spatially extended states to function as inner content bearers. Connectionist approaches helped us see beyond that vision, by identifying the content bearers as distributed patterns of activity. But it may take the full firepower of dynamic systems theory to reveal the rich and complex space of possible content bearers.

E. COGNITIVE INCREMENTALISM: THE BIG ISSUE

The work in artificial life (Chapter 6) and dynamic systems raises, in an especially acute form, a puzzle that we have already touched on several times. I think it is worthwhile, however, to now make this puzzle as explicit and prominent as possible.

The puzzle is this: What, in general, is the relation between the strategies used to solve basic problems of perception and action and those used to solve more abstract or higher level problems? Can the capacity to solve more intuitively "cognitive" problems (such as planning next year's vacation, thinking about absent friends, and designing a particle accelerator) be understood in essentially the same terms as the capacity to follow walls, to coordinate finger motions, to generate rhythmic stepping, and so on? Certainly, much of the recent literature on "embodied cognition" seems committed to a notion that I am calling "cognitive incrementalism." This is the idea that you do indeed get full-blown, human cognition by gradually adding "bells and whistles" to basic (embodied, embedded) strategies of relating to the present at hand. It is just such a principle of continu-

[11]For a similar argument, see Keijzer and Bem (1996).

[12]Just such a union is pursued in Crutchfield and Mitchell (1995) and in Mitchell et al. (1994). van Gelder's own notion of "revisionary representationalism" and his discussion of decision field theory (van Gelder, 1995, p. 359–363) show that he is open to the idea of such a union.

ity that prompts Thelen and Smith, for example, to comment that "there is in principle no difference between the processes engendering walking, reaching, and looking for hidden objects and those resulting in mathematics and poetry—cognition [is] seamless and dynamic" (Thelen and Smith, 1994, p. xxiii). Much depends, of course, on what we are here to understand by the phrase "no difference between." For in many interesting instances (see also Section B) we can discern both a kind of (often structural) continuity alongside some quite radical functional discontinuity. As a result, some cognitive functions may depend *not* on the tweaking of basic sensorimotor processing, but on the development of relatively (functionally) independent and (functionally) novel kinds of neural processes.

A case in point looks to be the "two visual systems" hypothesis of Milner and Goodale (1995). Milner and Goodale's claim, very (very!) briefly is that on-line visuomotor action is guided by neural resources that are quite fundamentally distinct (see Box 7.2) from those used to support conscious visual experience, off-line visual reasoning, and visually based categorization and verbal report. The latter complex of activities depends, it is argued, on a ventral processing stream and the former on a largely independent dorsal stream. Milner and Goodale's (admittedly quite contentious) explanation thus invokes a quite radical *dissociation* of codings-for-on-line action and for off-line reason and imagination. Here, then, is one very concrete case in which we seem to confront not a simple incremental process in which off-line reason exploits the very same basic mechanisms as on-line action guidance, but something more dramatic and different: a case, perhaps, in which nature adds functionality by developing whole new ways of processing and exploiting sensory input.

Notice that the Milner and Goodale story (unlike the example in Section B) does *not* depict reflective thought as simply the "off-line" use of strategies and encodings developed to promote fluent action in the here and now. Instead, it depicts nature as building (though doubtless out of old parts!) a *new kind* of cognitive machinery, allowing certain animals to categorize and comprehend their world in novel ways that are geared to the conceptualization of sensory input via the extraction of viewer-independent information (concerning object shape, identity, function, and so on). Such modes of encoding format, package and poise sensory information for use in conceptual thoughts and reason, and create what Milner and Goodale (1998, p. 4) suggestively call a system for "insight, hindsight and foresight about the visual world."

It is not my purpose, here, to attempt to fully describe, or critically assess this proposal (see Clark, 1999a, for an attempt). Rather, I invoke it merely to illustrate the empirical uncertainties hereabouts. It may indeed be—as Thelen, Smith, and others suggest—that the neural mechanisms of higher thought and reason are fully continuous with mechanisms of on-line action control. But it may be quite otherwise. Most likely, what we confront is a subtle and complex mixture of strategies, in which new kinds of information-processing routine peaceably coexist with, and

Box 7.2

Vision for Action versus Vision for Perception?

Milner and Goodale's (1995) controversial suggestion, briefly discussed in the text, is that the neural systems underlying visually guided action (such as reaching) are quite distinct from those underlying conscious visual recognition, categorization, experience, and imagination. A suggestive demonstration involves the so-called Tichener circles illusion (see Figure 7.3)—a case of illusory size distortions in which we regularly misjudge the sizes of the central discs. In the topmost drawing, the two central discs are (in fact) equal in size, whereas in the lower drawing they are different in size. The surrounding rings of large and small circles, in each case, lead us to perceptually misrepresent the actual size of the central discs, seeing them as different when they are the same (top case) and as the same when they are different (bottom case).

Perceptual experience here delivers a content that plainly misrepresents the actual size of the center discs. But there is a twist. Aglioti, Goodale, and Desouza (1995) set up a physical version of the illusion using thin poker chips as the discs, and then asked subjects to "pick up the target disc on the left if the two discs appeared equal in size and to pick up the one on the right if they appeared different in size" (Milner and Goodale, 1995, p. 167). The surprising result was that even when subjects were unaware of—but clearly subject to—the illusion, their motor control systems produced a precisely fitted grip with a finger-thumb aperture perfectly suited to the *actual* (non-illusory) size of the disc. This aperture was not arrived at by touching and adjusting, but was instead the direct result of the visual input. Yet, to repeat, it reflected not the illusory disc size given in the subject's visual experience, but the actual size. In short:

> Grip size was determined entirely by the true size of the target disc [and] the very act by means of which subjects indicated their susceptibility to the visual illusion (that is, picking up one of two target circles) was itself uninfluenced by the illusion. (Milner and Goodale, 1995, p. 168)

This is, indeed, a somewhat startling result. It suggests, to Milner and Goodale, that the processing underlying visual awareness may be operating quite independently of that underlying the visual control of action. Nor is this suggestion warranted only by the (interesting but perhaps somewhat marginal) case of these visual illusions. The general idea of a dissociation between systems for visual awareness and systems for visuomotor action is also suggested by anatomical data and data from brain-damaged patients. The patient DF, for example, suffers from ventral stream lesions and cannot visually identify objects or visually discriminate shapes. Nonetheless, she is able to pick up these very same objects—that she cannot visually identify—using fluent, well-oriented precision grips. Others, by contrast, suffer dorsal stream lesions and "have little trouble seeing [i.e., identifying objects in a visual scene] but a lot of trouble reaching for objects they can see. It is as though they cannot use the spatial information inherent in any visual scene" (Gazzaniga, 1998, p. 109).

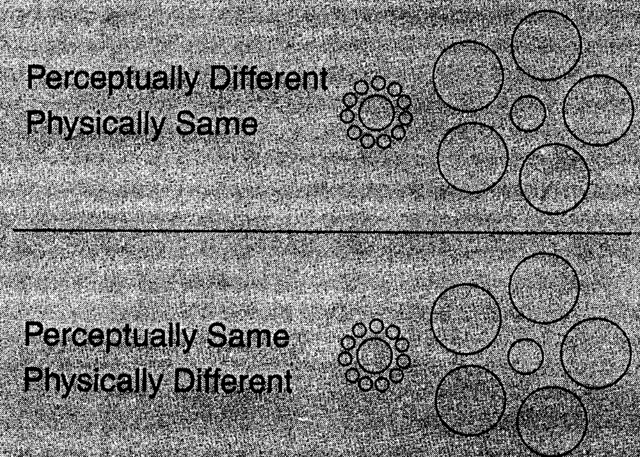

Figure 7.3 Diagram showing the Titchener circles illusion. In the top figure the two central discs are of the same actual size, but appear different; in the bottom figure, the disc surrounded by an annulus of large circles has been made somewhat larger in size in order to appear approximately equal in size to the other central disc. (From Milner and Goodale, 1995. By permission.)

at times exploit and coopt, more primitive systems (For some fascinating conjecture about the possible shape of such an interplay, see Damasio, 1999).

In sum, we must treat the doctrine of cognitive incrementalism with great caution. It is a doctrine that is both insufficiently precise (concerning what is to count as continuity, incremental change, etc.) and empirically insecure. Attention to the shape of nature's solution to basic problems of real-time response and sensorimotor coordination will surely teach us a lot. Whether it will teach us enough to understand mindfulness itself is still unknown.

7.3 Suggested Readings

For *accessible introductions to dynamical systems theory*, try R. Abraham and C. Shaw, *Dynamics—The Geometry of Behavior* (Redwood, CA: Addison Wesley, 1992); the chapter by A. Norton, "Dynamics: An introduction." In R. Port and T. van Gelder (eds.), *Mind as Motion* (Cambridge, MA: MIT Press, 1995); or (perhaps best of all for philosophers and cognitive scientists) Chapters 1–3 of J. A. Scott Kelso, *Dynamic Patterns: The Self-organization of Brian and Behavior* (Cambridge, MA: MIT Press, 1995, Chapters 1–3), which also contains descriptions of the work on *rhythmic finger motion*.

For the work on *infant stepping*, see E. Thelen and L. Smith, *A Dynamic Systems Approach to the Development of Cognition and Action* (Cambridge, MA: MIT Press, 1994), and critical discussion in A. Clark, "The dynamical challenge." *Cognitive Science*, 21(4), 461–481, 1997.

For the *Watt governor argument*, see T. van Gelder, "What might cognition be if not computation?" *Journal of Philosophy*, 92(7), 345–381, 1995, and critical discussion in A. Clark, "Time and mind." *Journal of Philosophy*, XCV(7), 354–376, 1998.

A good window on the *debate over internal representations* is provided by looking at on the one hand, A. Clark and J. Toribio, "Doing without representing?" *Synthese*, 101, 401–431, 1994, and on the other, F. Keijzer, "Doing without representations which specify what to do." *Philosophical Psychology*, 11(3), 267–302, 1998. The latter is a philosophically astute defense of a fairly radical dynamicist position, whereas the former is somewhat more skeptical.

The collection, by R. Port and T. van Gelder (eds.), *Mind as Motion* (Cambridge, MA: MIT Press, 1995) contains a number of interesting and provocative papers. I especially recommend the introduction "It's about time," by van Gelder and Port, "Language as a dynamical system," by Jeff Elman (a nice blend of connectionism and dynamics), and the robotics-oriented paper by R. Beer, "Computational-dynamical languages for autonomous agents."

For further discussion of the broad notion of *cognitive incrementalism*, see J. Fodor, *In Critical Condition*, (Cambridge, MA: MIT Press, 1998, Chapter 17, pp. 203–214).