Human Movement Science 10 (1991) 447–483
North-Holland

Self-organization in connectionist models: Associative memory, dissipative structures, and thermodynamic law *

Gavan Lintern

University of Illinois at Urbana-Champaign, Urbana, USA

Peter N. Kugler

Radford University, Radford, USA

Abstract

Lintern, G. and P.N. Kugler, 1991. Self-organization in connectionist models: Associative memory, dissipative structures, and thermodynamic law. Human Movement Science 10, 447–483.

In the development of connectionist models it is popular to rely on the concept of self-organization and to employ analogies from thermodynamics. Here we review some aspects of self-organization and thermodynamic law. We conclude that they do, indeed, have much to offer the modeling of human action. However, we further conclude that connectionists have failed to exploit the full potential of the properties inherent in a thermodynamic model of self-organization. Their use of self-organization lacks the imperatives of physical theorists or biologists who have written extensively on the topic. The use of computational temperature as an ordering principle for associative memory is analyzed. The more common approach in connectionism, to seek order through cooling, has less potential to explain the emergence of new behavioral properties than an approach that seeks order through heating. Thermodynamics as a source of analogies is also seen as limiting and we question the value of analogy as a basis for a scientific endeavor. An appeal to the constructive role of the Second Law as it operates on open systems can account for important features of organized activity. In this view the Second Law does not offer analogies; it is a law that describes the causal basis of human action.

* Preparation of this paper was supported by contract MDA 903-86-C-0169 to the Army Research Institute. Dr. Michael Drillings is the technical monitor. Aaron Contorer and Kim Vicente reviewed earlier versions of this paper.

Requests for reprints should be sent to G. Lintern, University of Illinois, Institute of Aviation, Q5, Aviation Research Laboratory, 1 Airport Road, Savoy, IL 61874, USA.

Introduction

Following a long period in which little interest was expressed in mechanistic accounts of motor control, an emerging concern with theories of cognition encouraged the development of motor theories that drew heavily on information processing and representational concepts (Adams 1971; Pew 1970; Schmidt 1975). These attempts at explanation were presumably driven by a preference for the deductive scientific strategy that a theory of mechanism allows (Casti 1989) but they raise a concern about the problem of infinite regress in that they assume a mechanism that incorporates the detail to be explained (Kugler 1986). In addition, serious questions can be raised about their biological plausibility.

In accordance with Casti (1989) we take the position that a theory of mechanism is desirable or even essential for the study of human behavior. Our interest is in the development of a theory of human action for which the origin and evolution is biologically plausible. A specific concern is with nomination of the semantic primitives. How can a mechanism for intelligent or organized behavior evolve without an a priori description of that behavior being provided by an external agent? Biological plausibility and nomination of semantic primitives are, we believe, foundational issues for the modeling of a mechanism that can support human action.

The emerging discipline of Parallel Distributed Processing offers the possibility of a solution. Proponents claim that their models are biologically plausible and that semantic primitives are not inserted into the system by the theorist (Smolensky 1988). It has been claimed that the approach constitutes a paradigm shift (Schneider 1987), and it is one that might be as effectively applied to human action as to human cognition. In particular, we view organized human activity in a perceptual-motor workspace as supported by a style of cognition which is conceptually no different to the style of cognition generally envisioned within the connectionist literature. From that perspective, those interested in mechanisms underlying the organization of human movement should profit from understanding what is going on. In this paper we review fundamental concepts that drive some of the connectionist work and evaluate their use by theorists involved in connectionist modeling. In addition, we consider the more general value of those concepts in the study of human behavior.

The problem

One central theme to Parallel Distributed Processing (PDP) is that of self-organization. Rosenblatt's 'vision of the human information processing system as a dynamic, interactive, self-organizing system lies at the core of the PDP approach' (McClelland et al. 1987: 42; in reference to Rosenblatt 1959, 1962). Self-organizing systems are necessarily dynamic and interactive. Thus we are left with the issue of what it means to be self-organizing. A second significant theme running through PDP is an appeal to thermodynamic law as a source of analogies for cognitive process. What is to be gained by delving into thermodynamic law that is not already well enough covered in other contemporary accounts? Furthermore, is it satisfactory for a new scientific thrust to be based on analogy?

In this paper we will explore several issues. The first is the concern within the connectionist literature for micro and macro descriptions and the mapping between them. The second is the use of thermodynamic law in cognitive theorizing. We will contrast cooling versus heating as principles for the emergence of order from disorder within the context of self-organization and will explore the notion of an associative memory. Finally, we will contrast the roles of analogy and law as a basis for scientific progress.

Levels of description

The goal of PDP is to offer 'computationally sufficient and psychologically accurate mechanistic accounts of the phenomena of human cognition' (McClelland et al. 1987: 11; Ballard 1986: 67). To that end connectionists are primarily concerned with the micro (mechanistic) structure of cognition although they recognize that macro (phenomenal) processes influence behavior. In particular, there is a regard for emergent phenomena that could never be predicted or understood from an isolated description of the micro structure but can be understood in terms of interactions within the micro structure (Rumelhart and McClelland 1987: 128). In the PDP view, entities referred to at the macro level of behavior are approximate descriptions of emergent properties of the micro structure (McClelland et al. 1987: 12).

Forms of description

It is to be hoped that the meanings of these statements have not been distorted out of context, because we now wish to make some subtle but important distinctions about the nature of description. At a single level, there can be *alternate* forms of description; more detailed, more abstract, or merely based on a different language or concepts. There can also be alternate descriptions between levels. Classical reductionism supposes that macro phenomena can be described in more detail at a micro level and if the micro concepts are available there is little to be gained from description at the macro level. According to this view the study of macro behavior results in an alternate description of temporary value to be replaced by neurological, and possibly physical and chemical descriptions when enough is known at the micro level.

A contrasting view is that descriptions at the micro and macro levels are *complementary* (Pattee 1979); that is, they mutually enhance each other but are not fundamentally reducible one to the other. Accurate descriptions at both macro and micro levels plus knowledge of how the micro phenomena map into the macro phenomena (and, in a self-organizing system, how macro and micro states interact) are essential to understanding mechanisms that result in organization at the macro scale. From this perspective the PDP enterprise seems well motivated. Nevertheless, clarification has been offered here because a consideration of macro states does not always appear to impose any significant constraints on PDP theorists (e.g., Smolensky 1987) and there is the occasional implication that an understanding of interactions within the micro structure is sufficient for a full understanding of cognition (Rumelhart and McClelland 1987: 128).

Between-scale mappings

The view developed in this paper is that a consideration of both micro and macro levels is crucial. They are complementary and irreducible descriptions. Neither can be deemed the more fundamental or primary and neither can serve as an approximate description of the other. Furthermore, an important key to furthering understanding of a (cognitive) mechanism lies in exploration of the mapping between the two levels (fig. 1). Lucid and accurate descriptions are required at both levels before the mapping problem can be solved. It would seem,

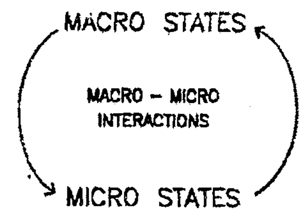

Fig. 1. Natural phenomena may be described at multiple levels or scales. In the classical reductionist view these are alternate descriptions with the micro description constituting the scientifically more significant level. From the self-organizational perspective descriptions at two or more levels and knowledge of interactions between levels are essential for full understanding.

nevertheless, that there are useful (although as yet partial) descriptions at both the macro level (human behavior) and the micro level (neurology, biochemistry, etc.). In many developments of PDP models connectionists have embarked on the important task of exploring the mapping between the two levels.

Micro and macro descriptions of brain and behavior

Norman (1987: 534) observes that PDP models treat psychological data within constraints imposed by neurological data. This approach is consistent with Crick's (1979) view that an explanation of human behavior cannot be derived either from psychology on its own or from the neurosciences in isolation from psychology. Knowledge from both is essential with the further need to understand how the micro states and processes studied by the neurosciences are linked to the macro behavioral states studied in psychology (fig. 2). For Crick, communication theory (which appears to encompass notions of how neuronal

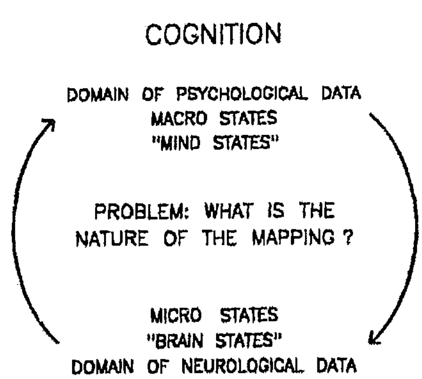

Fig. 2. An important key to understanding cognition lies in understanding the mapping between macro and micro levels.

activity is transformed into information and how that information is stored and retrieved) provides the appropriate link, while for Norman, computational processes fulfill that requirement.

Crick (1979) has summarized several important constraints for a theory of human behavior. One is the nature of the environment and our interaction with it; that is the macro description of the events under study. Another set of constraints is related to structure and processes of the central nervous system. The multitude of neurons (approximately 10^{11}) and connections between them (approximately 10^{15}) together with their relatively slow action (in the order of milliseconds) would seem to be essential considerations. A third set of constraints relate to the global characteristics of the neural system; the fact that there are precise connections between neurons in some parts of the central nervous system and that there are discrete areas in the cortex in which activity is of the distributed nature of an associative net. In addition, it is necessary to avoid conceptions that rely on the intelligence of single neurons or a homunculus for explaining the emergence of organized (cognitive) activity.

In modeling the micro structure of human behavior it would be most desirable to rely entirely on neuroscience and biochemistry for definition of structures and processes. Unfortunately, there are huge gaps in our knowledge of the anatomy and physiology of the brain and there are enormous technical and ethical obstacles to gathering the desired information (Crick and Asanuma 1987). The obstacles are so significant that any comprehensive account of the micro structure of cognition must postulate a large number of hypothetical structures or processes that have no definite support from neuroscience.

While PDP models are not intended to describe the detailed neural implementation of behavior (Rumelhart and McClelland 1987: 138) they are said to be neurally inspired (McClelland et al. 1987: 11; Norman 1987: 535) at least to the extent of the multitude of many-to-one and one-to-many connections. A pervading attitude within the PDP enterprise appears to be at least approximately consistent with Crick's view: PDP models tend to be loosely constrained by some established facts of neuroscience while PDP theorists take license to supplement the facts with additional processes or structures that have no neurological support. The challenge is to motivate from first principles the selection of these additional processes. Without first principles to bootstrap the selection process, proposed processes can be considered little more then Kiplinese 'just so stories'. For this paper the issue we will focus on is associative memory: what are its characteristics, what are the essential mechanisms, and what first principles might we appeal to in constructing an account of it?

Associative memory

The essential components of an abstract, computational model of the brain can be described in terms of synaptic connections within neural networks (Ballard 1986: 67) or an associative memory in which patterns of activity represent modal states for perceptual and cognitive behavior (Baird 1986). An associative memory is a network of interconnected elements that produce distinctive patterns in parallel outputs from patterns of excitation in parallel inputs. The macroscopic pattern of a modal state reflects the mass action of microscopic neural events throughout the associative memory. Knowledge resides in the pattern of neural connectivity rather than in any single neuron or group of neurons or in any feature processing subsystems.

The parallelism of the system is important from the timing point of view. If transitions between cognitive (macroscopic) states are as frequent as two or three per second (McClelland et al. 1987: 12; Ballard 1986: 67) and the transitions between neuronal (microscopic) states require several milliseconds, serial processing will be too sluggish. Thus, the concept is of an associative network that responds to different inputs by settling into distinctive but distributed states. The strength of the connectivity between elements and the subtle interactions between activation and inhibition rules determine the reliability of the settling state in response to specific inputs. Presumably cognition is based on the assembly and disassembly of macroscopic patterns of activity at rates that range at different levels of organization from a few per second to a few per lifetime.

In a neural network there are probably a countless number of potential modal states. Through learning or experience, some of these will become preferred states (solutions) in that they are the ones most likely to be activated during normal activity. We assume that a large number of latent preferred states can coexist within a neural network although that large number will be considerably smaller than the countless number of pre-existing potentialities. We further assume that only one modal state is active at any one time, and that a neural network progresses through cycles of activation → deactivation → activation... and so on. These assumptions about associative memory appear to be generally consistent with prevailing connectionist views. Crucial issues for this view are how specific modal states acquire their preferred status and, once several are established, how a state might be selected and re-selected for activation.

Simulated annealing: Crystal formation as an optimum solution

One regularity accounted for by Thermodynamic Law is the tendency, noted by McClelland and Rumelhart (1988: 70), 'for all physical systems to evolve from highly energetic states to states of minimal energy'. This is an approximate characterization of Boltzman's ordering principle which specifies that the direction of natural change within an isolated system is towards a state of maximum homogeneity. It is the principle behind the order-from-cooling analogy, as exemplified in crystal formation, which has been inspirational for PDP modelers as a means of showing how a neural network might settle into a modal state or how it might achieve an optimum solution (e.g., Hinton and Sejnowski, 1987; Hopfield and Tank, 1986; Smolensky 1987).

For the crystal formation analogy it is important to appreciate that a reduction in thermal agitation reduces molecular kinetic energy. Some of that kinetic energy is lost to the surround and the lower level of thermal agitation also permits some potential energy, which is stored within molecular bonds, to be given up and lost to the surround. The most orderly state is one that ends up at the lowest (minimum) potential energy state. Thus, a pure crystalline substance (regular molecular pattern) is in a minimum potential energy state (lattice configuration) and can be said to offer an ideal solution (symmetric lattice). An impure crystalline substance is one in which the alignment of the lattice pattern is not identical throughout the crystal. Potential energy is low; it is minimized in local regions, but it is not minimized globally. This could be referred to as a nonoptimum solution. In some circumstances it might be viewed as a near-optimum solution and in others, as an error.

A pure crystal can be obtained by first heating the substance to break the crystalline bonds and then applying an appropriate cooling schedule. With careful cooling (i.e., slow reductions in temperature over the critical range at which the crystal forms), a portion of the substance will, by chance, crystallize first, and adjacent portions will take on the same structure. That structure will thereby permeate through the substance. This type of careful cooling is termed 'annealing'. In contrast, rapid cooling (or 'quenching') will precipitate crystal formation in different portions of the substance at the same time. The orientation of the crystalline structure will not be identical throughout. The meeting of different orientations in the structure constitutes an impurity or an intersection at which energy is not minimum. Thus, potential energy of the crystal lattice is minimized locally in most parts of the crystal but it is not minimized globally.

Computational temperature

It is this capacity of molecular configurations to locate globally optimum solutions in configuration spaces having many thousands of local minima that has guided the development of some PDP models. Such systems with many local solutions competing with a global solution are referred to as frustrated systems. Computational temperature (fig. 3), as an analog to thermodynamic temperature, is reduced slowly over the critical range so that the system settles into a 'harmonious' or 'minimum energy state' (Hopfield and Tank 1986; Smolensky 1987). This minimum energy state is also characterized as an attractor state (McClelland and Rumelhart 1988: 70). Nevertheless, it should be noted that this system does not learn; there is no residue from the

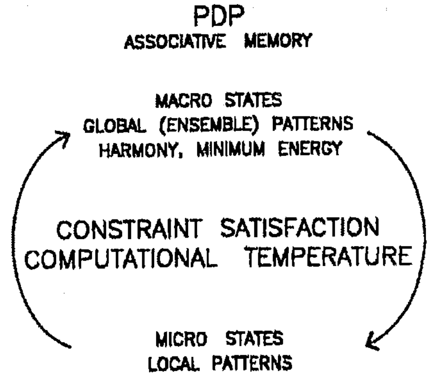

Fig. 3. Many PDP models account for the mapping between macro and micro levels by constraint satisfaction.

achievement of an optimum state that will help the system locate that same state more easily on a subsequent heating and cooling cycle.

Constraint satisfaction

One welcome feature of the PDP approach is that some models have sought to deal with the issue of how a preferred state can be selected for activation (recognized) without the need for an a priori global or macro-state reference (i.e., an internal representation, a set point, or homunculus). As is evident from the crystal formation analogy, one approach taken is to minimize or maximize a natural global (macro) process within the system by ordering its elements (fig. 3). Smolensky (1987) has chosen to maximize harmony within the network (i.e., to achieve a maximally self-consistent state), while Hopfield and Tank (1986) have chosen to minimize computational energy. There is no crucial distinction to be made between maximization and minimization approaches which can both be thought of as resulting in optimization via constraint satisfaction.

In essence, the problem is one of ensuring that a regular input can reliably produce the same activation state but one that differs from states produced by other inputs. A beginning state of random activity in which elements of the model are stochastically activated is assumed. The level of random activity is used to define a computational temperature. At a high computational temperature the random activity is so high that a regular input cannot exert any ordering influence. As computational temperature is lowered the level of random activity is also reduced and the regular input can start to exert its influence. The excitatory and inhibitory connections interact with the input to guide the system towards the desired state. As temperature is further lowered random activity ceases and the system becomes frozen in a final state.

The decrease in computational temperature must be scheduled carefully. A sudden decrease ('quenching') may locate part of the system in a stochastically determined minimum that is a locally (but not a globally) optimum solution.. Carefully scheduled cooling (simulated annealing) can result in the system settling into a global optimum. Under the appropriate schedule the regular or coherent input gradually establishes its influence over the disorganized activity of the system in a manner that results in convergence onto a globally defined optimum state (Metropolis et al. 1953).

Solutions as attractor states

Terminology from nonlinear dynamics is occasionally employed in the connectionist literature (e.g., Baird 1986; Hopfield and Tank 1986; Skarda and Freeman 1987). When a system settles on a solution it can be said to converge onto an attractor [1]. The initial conditions from which the system converges onto an attractor are within the basin of that attractor, and the entire set of initial conditions that leads to convergence on a given attractor constitutes the basin for that attractor. Multiple attractors coexist in some systems. The selection of an attractor state, although a deterministic function of the initial conditions, is often difficult to predict because of limitations in measurement of the initial conditions. In particular, the difference between initial conditions that are in the basin of one attractor versus another can be so small that it cannot (even in principle) be measured reliably and the system trajectory can be said to be infinitely sensitive to initial conditions.

Categorical perception

One intriguing aspect of this terminology for psychology is that it may be used as a description of categorical perception. Initial conditions in the form of stimulus information, that may be discriminated with the assistance of special instrumentation, are classified as identical by the unaided human perceptual process. Thus, the terminology of attractor states represents an alternate style of description for behavior that is often described in terms of templates, schemas, or internal models.

An analogy from Baird (1986) for categorical perception is that of a flexible buckling column with a deformable collar. A vertical column will remain straight with the addition of a downward force until a threshold is reached, when it will buckle or bow to an extent de-

[1] An attractor defines an invariant solution shared by multiple trajectories originating from different initial conditions. It is a global symmetry that relates local trajectories.

termined by the force. The column can buckle in any direction. The precise direction taken is determined by micro fluctuations in the column at the time that the vertical force exceeds the threshold required to initiate the buckle. The collar will be deformed by the buckle and more deformations can be created by releasing and reapplying the vertical force. Nevertheless, a buckle in the general direction of an existing deformation will not create a new impression but will, instead, be captured by the existing deformation. Where once behavior (the direction of buckling) was unconstrained it is now constrained to discrete states that depend on prior experience. A memory has been created and that memory may be described succinctly in terms of multiple latent attractors.

Prigogine's principle: Order through heating

The crystal formation analogy may be viewed as drawing on an order-through-cooling principle to explain emerging organization in a cognitive system. We now return to our earlier discussion of micro and macro descriptions so that we may contrast 'order through cooling' with 'order through heating'. Specifically, in many closed and open systems, flows of energy (closed systems) or of energy and matter (open systems) can create a spontaneous transition from one ordered state (or from an homogeneous state) to a new ordered state in a manner that is not consistent with Boltzman's ordering principle.

Self-organization in a closed system

The Rayleigh-Benard instability offers a classic example of a closed, self-organizing system in which order arises out of a heating process (Haken 1981; Berge et al. 1984). A homogeneous layer of thermally expansive fluid, if heated uniformly, develops a regular structure of thermal convection rolls with parallel, horizontal axes. As the fluid is heated from below, a vertical temperature gradient is created and the lower layers of liquid expand to become less dense. At small temperature gradients the tendency for the lower and lighter (less dense) portions of the liquid to be displaced by the upper, heavier (more dense) portions is resisted by the dynamic viscosity (friction) of the liquid. At this temperature only heat conduction occurs. Once the temperature gradient becomes sufficiently strong, heat conduction is replaced by convection of matter resulting in the formation of convection rolls. Some useful implications concerning the role of macro-micro interactions can be drawn from the order-through-heating principle as manifested in the Rayleigh-Benard instability.

A symmetry-breaking nonlinearity

The first is that the motion threshold of the temperature gradient is a nonlinearity produced by a competition between forces. Thermal forces compete with gravitational and viscous forces for control over the matter transport. The convection rolls (circular matter transports) emerge when the temperature gradient generates sufficient force to overcome the gravitational and viscous forces. In general terms, the system response to a gradual increase in temperature gradient is nonmonotonic. A symmetry break occurs at a threshold temperature gradient. One state (in this case a homogeneous or disordered one) transitions suddenly into a qualitatively different state. Prior to the symmetry break the system response is linear in that an increase in temperature gradient has a monotonic effect on thermal forces. The abrupt transition to a qualitatively different state represents a nonlinear transition which is then followed by another region of linear behavior. More complicated patterns may emerge when further symmetry breaks occur at higher temperature gradients.

Micro–macro action

A second insight to be drawn from the convection example is that a transition from homogeneity to structure (a symmetry break) is understood more clearly by distinguishing the micro from the macro states. At first, heating increases molecular thermal agitation in the micro states which, in turn, increases the kinetic energy exchanges between molecules (conduction) without changing the lattice configuration of the homogeneous macro state. It is these forces which lead to thermal expansion and thus to the reduced density of the heated liquid, that have an important role to play in breaking the symmetry of the homogeneous lattice configuration by initiating and then sustaining the molecular transports forming the convection rolls (the new macro state). The rolling motions are sustained by the temperature gradient that occurs because the liquid at the bottom of the vessel gains heat while the liquid at the top loses it to the environment (fig. 4).

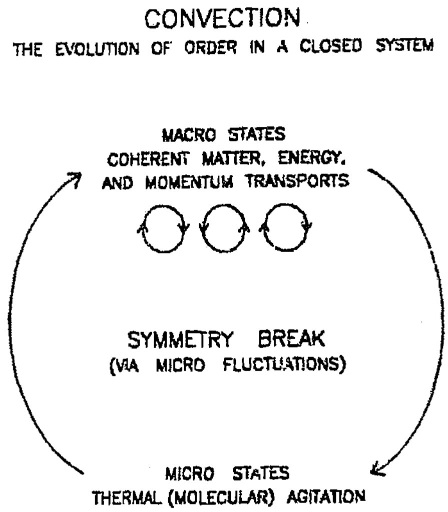

Fig. 4. Order in a closed system evolves through a spontaneous symmetry break initiated by micro fluctuations and sustained by competition between heating and cooling processes.

Macro selection via micro fluctuations

In a rectangular vessel the axes of the convection rolls are parallel to the shorter pair of sides. The direction of motion alternates between adjacent rolls but the motion of a specific roll may, when viewed in cross section, be clockwise or counterclockwise. The selection of a direction is made by a fluctuation at the time of transition from homogeneity to order. With the increase in thermal agitation portions of liquid are displaced in random fashion. At one moment, by a chance fluctuation, more will be displaced in a direction that favors one pattern over the other. At low temperature gradients such fluctuations are damped out. At the critical gradient a fluctuation generates a mini roll that enslaves other nearby elements in that direction of motion. The progressive enslavement moves quickly through the liquid to establish the convective transport pattern. This is a process that can be referred to as a spontaneous break in the symmetry of the homogeneous state.

Macro–micro action

Because the pattern has two equiprobable states the system is said to bifurcate at the transition from homogeneity to order. That transition point is known as the point of bifurcation (see fig. 5). The directional characteristic of the macro pattern is thus led by a stochastic fluctua-

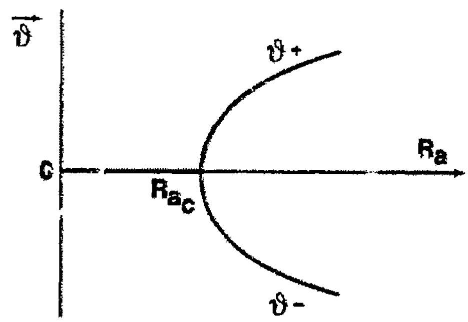

Fig. 5. The bifurcation diagram for the Rayleigh-Benard convection. R_{a_c} is the Rayleigh number (a dimensionless number) proportional to the ratio of the dynamic viscosity of the liquid and temperature gradient. ϑ is the velocity of a specific roll. In an experiment free of imperfections, the rolls engendered at R_{a_c} have equal probability of rotating in either direction. This is expressed by the existence of two branches denoted by ϑ_+ and ϑ_-.

tion at a micro level; a micro–macro action. However once selected, the macro pattern enslaves the micro motions via a macro–micro action. The macro organization ensures that one of the previously equiprobable patterns now dominates system behavior.

One feature of convection not found in crystal formation is the coexistence of multiple ideal solutions. However, like crystal formation, there is no residue that can be classified as memory and to achieve that we must turn to an example of an open system in which there are flows of matter as well as of energy.

Self-organization in open systems

The periodic assembling of a nest by a population of social insects using pheromone gradients provides an illustration of how generic symmetry breaking and selection mechanisms can function in a biological system. In summary of an account by Kugler and Turvey (1987; also see Deneubourge 1977; and Grasse 1959), African termites are known to construct nests from their waste deposits. After an initial random-deposit phase they commence construction of pillars. With some frequency, neighboring pillars can take on a mutual curvature towards a virtual midpoint so that an arch is constructed. This is followed by the construction of a dome over the supporting pillars and arches, and then a new construction cycle of pillars, arches, and a dome may begin (fig. 6). This structure becomes the nest in which the termites live and breed. The construction involves the coordination of

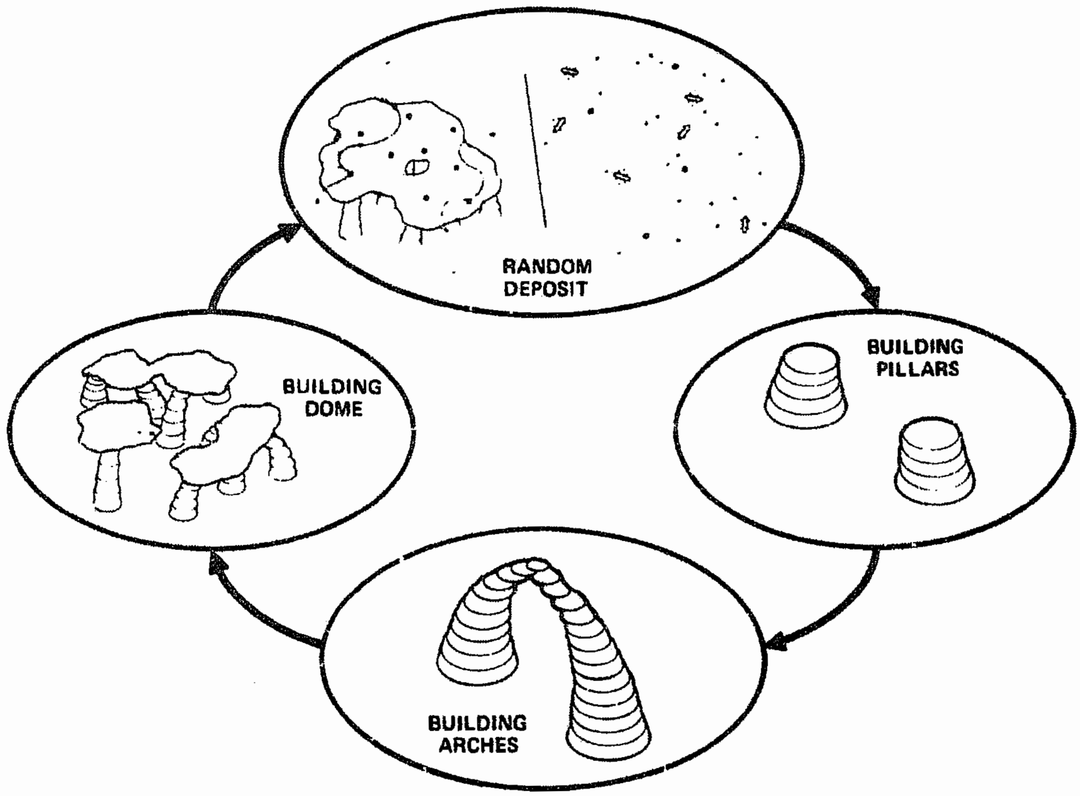

Fig. 6. A building cycle commencing with a random deposit phase and proceeding through pillar construction, arch construction, dome construction and a return to the random deposit, where the cycle may start again (adapted from Kugler and Turvey 1987).

several million essentially identical elements (the termites) without the benefit of rules or of a plan.

A symmetry-breaking nonlinearity

Construction starts with disorganized behavior in which termites fly through an area leaving behind waste deposits which contain a chemical pheromone that attracts other termites. Although the potency of the pheromone decays relatively quickly the frequency with which insects fly within their detection threshold of the pheromone field from an active site before it decays will, if a sufficient number of insects participate, become relatively high. Insects that fly within their perceptual threshold of an active site will change their flight trajectory to pass over that site and will deposit waste on it. Thus, a small number of preferred deposit sites can begin to emerge.

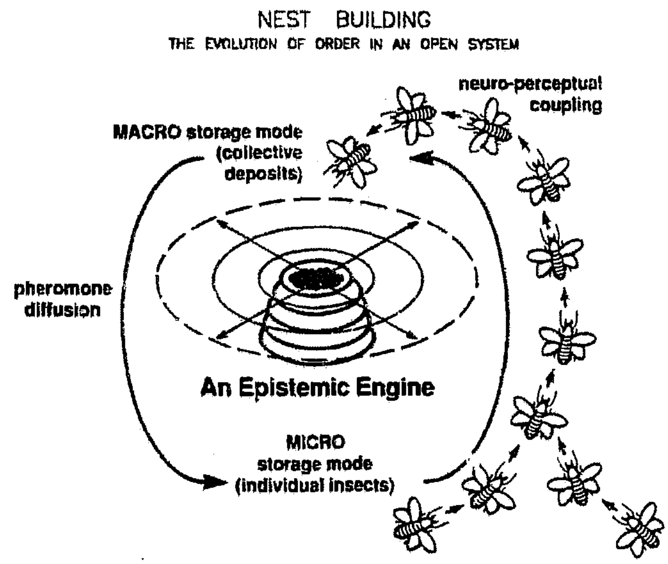

Fig. 7. The autocatalytic cycle of pillar construction by social insects (adapted from Kugler and Turvey 1987).

Micro–macro action

The locations of the preferred sites are selected through fluctuations in individual flight trajectories at the micro scale (of individual insects). Once a particular site is selected an autocatalytic reaction enslaves the ensemble of insects to form a macroscopic flight pattern organized by the pheromone diffusion field (fig. 7). The result is the formation of pillars of waste at preferred sites.

Emergence of a higher-order pattern

Arches begin to form when the pheromone fields of neighboring pillars strengthen (in terms of insects' perceptual thresholds) to the point that they begin to overlap. A mutual interaction between two sites can have a biasing effect that results in deposits accumulating more rapidly on the proximal sides of the two sites. Curvature in the pillars follows which may result in the two pillars eventually meeting to form an arch. Thus, another threshold is crossed and a new macro property emerges. When several pairs of pillars meet new field properties emerge which result in the construction of a roof.

Macro selection via micro fluctuations

As in the Rayleigh-Benard instability the specific macro pattern of sites that emerges from homogeneity is selected by virtue of stochastic fluctuations within the micro processes (a micro–macro interaction). A threshold is manifest that identifies a functional boundary separating a nonequilibrium relaxational dynamic from a 'far-from-equilibrium' self-organizing dynamic (cf. Prigogine and Stengers 1984). In the equilibrium condition insect behaviors are independent of one another. In the far-from-equilibrium condition the behaviors of individual insects are highly coordinated.

Macro–micro action

There is no requirement for the individual elements to communicate with each other although it might be said that information is fed back to individual elements from the macro level of organization via the pheromone field. The organization is by virtue of the mutual coherence of behavior; in particular the interactive dynamic of the termite sensory apparatus (micro) and the pheromone information field (macro). Through this process a macro pattern of organizing centers is assembled and this pattern, once established, effectively organizes the flight trajectories of individual insects (a macro–micro action).

Towards an open systems account of cognition

Dissipative structures

Patterns created from the changing balance of forces as induced by flows of energy (or of energy and matter) are known as dissipative structures (Prigogine and Stengers, 1984). Essentially, dissipative structures are new forms of order that emerge via state transitions when energy flows exceed the dissipative capacity of an existing structure. They emerge as a result of the instabilities produced in existing structures by those high rates of energy flow. The continued viability of a dissipative structure is maintained by these flows through the system with one causal influence residing in the forces that dissipate some of the energy into activity at the micro scale.

The continuous flow of energy is critical. If termites were constrained from flying through an area for a time the pheromone fields would decay and the pattern of organizing centers would be disassembled. Similarly, if the heat source is removed from beneath a thermally expansive liquid the convection rolls will decay and the system will return to the homogeneous state of stable equilibrium.

Closed versus open systems: Flows of matter and creation of memory

There is, however, an important distinction to be made between the convection and nest construction systems. In contrast to a closed system (e.g., the Rayleigh-Benard instability), an open system (e.g., nest construction), does not return to the homogeneous state at cessation of the energy (and matter) flows. There is a physical instantiation of the organization (the nest structure), which results specifically because of the open (versus the closed) nature of the system. This physical instantiation (which will eventually be torn down by Second Law processes) continues to constrain the flight trajectories of the insects as they fly through the nests and might be viewed as a memory of the nest building activity. Bertalanffy (1975) has argued for an open-systems account of biological processes (including cognition) and here we extend his arguments by observing that a principled account for the origin of symbolic, rate-independent constraints normally characterized as memory can be found in the dynamic, rate-dependent (Second Law) processes of open systems.

Self-organization: A nonrepresentational account

Self-organization is characterized by transitions to new states of order in the absence of any a priori material embodiment (occupancy of physical degrees of freedom) that specifies or represents a set point, representation, template, or schema in the medium from which the pattern is constructed or in the input to the system; that is the heat flow or the temperature gradient in convection or the transport of energy and matter by termites in the construction of nests (contrast with fig. 8). Furthermore, there are no special purpose elements. Haken (1981) attributes the emergence of structure to an order parameter; in essence an activity that enslaves other activity [2]. There are, however, constraints that may be viewed as control parameters; for example some characteristics of a convection pattern are determined by the shape of the holding vessel and perceptual thresholds determine the value of the order parameter that induces autocatalytic amplification of

[2] A self-evident example of an order parameter at work is found in an avalanche. A rolling boulder imparts its motion to other boulders that join the system, i.e., the motion of one boulder enslaves the motion of others.

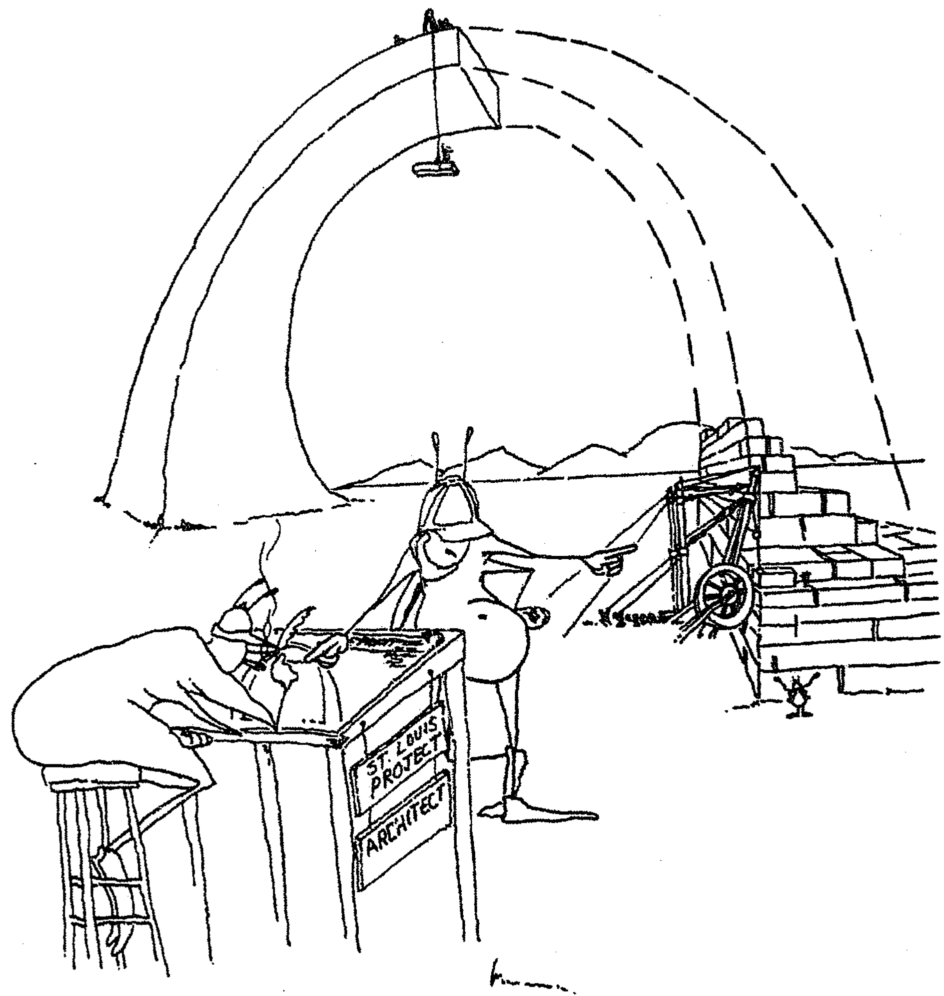

Fig. 8. Artifactual-machine solution for the construction of an arch. The long range correlation required for the cooperative building of the arch is accomplished by a small scale blue-print (plan, frame, schema, etc.) of the large scale project. The actions of the individual workers are constrained according to restrictions specified in the blueprint. External regularity is intimately tied to internal regularity of the blueprint. Of particular importance is the fact that the blueprint must exist prior to and independent of the actual construction. In general it takes an agent (i.e., an architect) more complicated than the phenomenon being explained to account for the origin of the blueprint (adapted from Kugler 1986).

fluctuations in termite nest construction. Nonlinearities (thresholds, switches, hysteresis, damping, inertia, saturation) abound in self-organizing systems and are influential constraints on the emergence of order and on its final state.

Although some characteristics of an emergent order are unpredictable (for example, the particular sites selected for termite nest construction) the forms that emerge can be recognized as characteristic of the system. Nevertheless not anything is possible. The emergent properties remain bounded within a range of possibilities. Termites may build pillar- or arch-like structures but they may not build representations of artifacts such as the lattice work found in the Eiffel Tower or the detailed features of the Statue of Liberty. This feature of open systems may also be seen as characteristic of human cognitive behavior; essentially unpredictable in detail but bounded within a range of identifiable patterns.

Generic mechanisms for self-organization

A comprehensive explanation of the emergence of structure from homogeneity (or the transition from one ordered state to another) requires a macro description of the structure (the emergent properties), a micro description of elemental activity, and some consideration of the mapping (or the interaction) between the micro and the macro states. In particular, appreciation of fluctuations in the micro structure is essential for understanding a bifurcation (i.e., selection). The emergence of new macro properties in a self-organizing system can be said to result from nonlinear transitions induced by amplification of fluctuations in the micro structure when the system is forced far from equilibrium (Prigogine and Stengers 1984: ch. 6). Organization emerges in the absence of any prior representation of that pattern either in the input, the medium, or rules that map micro to macro states and, once established, that organization serves to constrain system behavior. Again, explanation is aided by descriptions of both macro and micro states and an understanding of how the interaction of cooperating and competing forces can lead to a nonlinearity which, in turn, leads to a creation of new forms not specified in the micro structure.

Emergence of order: Cooling versus heating

The statement by McClelland and Rumelhart (1987) that all physical systems tend to a minimum energy state is generally correct but misleading in the sense that, as is evident from a consideration of dissipative structures, local fluctuations can produce local increases in energy that create new states of order. The states of unstable equilibrium that are achieved in the self-organizing systems described above might also be viewed as solutions in the sense of the term as used in Boltzman, Harmony, and simulated annealing models. There is, however, a distinction to be drawn. One type of solution is achieved through heating and the other is achieved through cooling.

On the surface, the distinction between the two processes may not seem to amount to much. Carefully scheduled cooling leads to a singular and optimum state (a state of stable equilibrium) with a high probability. That state can be disassembled and reassembled by reheating the system and then allowing it to resettle by cooling with the possibility that it will settle into a different state. Heating can also lead to an ordered state (one of unstable equilibrium) which can be disassembled and then reassembled first by cooling and then by reheating. The cooling and reheating may allow the system to settle onto a new state via inverse bifurcation.

There are, however, some noteworthy differences. For systems that organize by increasing energy flows there appear to be solutions within solutions. This is found in the Rayleigh-Benard convection in which new patterns emerge as the temperature gradient is increased, eventually giving way to chaos. It is also found in termite nest construction where a first solution is elaborated into a new and distinctive one. In particular, this example suggests that the elaborated solution retains much of the character of the first but also has additional distinctive features. It might be characterized as a higher-order solution or, in Gibson's (1979) terms, a higher-order invariant. In some systems the elaboration of solutions can proceed through a considerable number of bifurcations before there is a transition to chaos (e.g., predator–prey relationships as modeled by the logistic equation, May 1976).

There is another important objection to applying the sudden emergence of orderly structure in the phase transitions that accompany cooling to the phenomena of cognitive processes. As noted by Haken (1981: 41–42) 'life processes slow as temperature drops; in fact they stop completely at very low temperatures. Living organisms are kept alive by a constant supply of energy and matter which they take up and process. The more highly developed creatures (i.e., the warm-blooded animals) are not even in thermal equilibrium with their surroundings.' Haken concludes that life processes cannot be based on a cooling principle. The solutions achieved by heating are maintained only by a continuous flow of energy (a closed system) or a continuous flow of energy and matter (an open system). In this regard the heating principle is more consistent with the realities of biological function.

In the next few pages we deal with two parallel distributed models that incorporate some self-organizing features. The first is consistent with a closed system account and is modeled with cellular automata (Langton 1986). The second is consistent with an open system account and is based on EEG data collected from the olfactory bulb of the rabbit (Baird 1986).

Artificial life: Cognition in a closed system?

To model the emergence of structure with increasing temperature it would be necessary to develop a system in which low temperatures lead to macroscopic homogeneity and high temperatures produce macroscopic chaos. Interesting properties would have to emerge in the intermediate temperature ranges. To be fully consistent with the arguments outlined above, those interesting macro properties should be led by fluctuations in the micro states but should otherwise be self-organizing. An autocatalytic effect in which micro fluctuations are amplified into self-sustaining macro properties is likely to provide the mechanism.

Such self-organizing patterns can be found in systems of cellular automata. Cellular automata are mathematical models for complex natural systems in which there are local interactions between large numbers of simple identical components (Wolfram 1984). These systems, which are best known through Conway's game of Life (Atkins 1984), are essentially simulations in which the cells of a matrix can take on one of a finite number of states over successive cycles of the simulation. Matrices are generally one-dimensional ($1 \times n$) or two-dimensional ($n \times n$). A typical two-dimensional automaton may have 64×64 cells with each cell capable of taking on one of eight possible states.

Other important elements of the simulation are a quiescent state, a neighborhood, and a transition function. One of the possible cell states is defined as a quiescent state to represent nonactivity. A neighborhood is a pattern of nearby cells that can be affected by an active cell and will generally include the active cell itself. The transition function specifies how an active cell affects its neighborhood (i.e., what type of states it generates in its surroundings for the next cycle).

Self organization in cellular automata

A cellular automaton may be viewed as a massively parallel system with local connections. A simulation is started with a seed in which some of the cells are switched into active states. Because no pattern is specified by the transition function any new pattern occurring in successive iterations must be an emergent property. The system is deterministic in that the same initial conditions and transition function will produce the same behavior. The emergence of a structure or pattern will constitute a solution in the same sense that the settling of a PDP network constitutes a solution.

Langton's (1986) investigations of artificial life with cellular automata are of particular interest. He specified a parameter, Γ, which determined the relative level of neighborhood activity generated on successive cycles by an active cell. The Γ parameter can be viewed as a strength determinant for an autocatalytic effect that generates new activity from existing activity. That parameter might also be viewed as a measure of system temperature.

Order through heating

Langton started his simulations with a randomly generated transition function constrained by Γ, and a randomly generated seed. For a Γ of zero the activity in the next cycle must, by definition, collapse onto the quiescent state. For a non-zero Γ, activity continued though succeeding cycles, but for small Γ, it quickly collapsed onto the quiescent state. Beyond some threshold, activity continued in a self-sustaining mode. As Γ was increased from that threshold the generated activity progressed through the stages of emergent fixed or propagating single states, several species of emergent periodic propagating structures that met to interactively generate new fixed or periodic structures, and finally chaos. In Langton's words, 'For low Γ temperatures we observe precipitate-like behavior where everything is stable and nothing changes, while for high temperatures we observe the behavior of a hot gas where everything changes and nothing is stable. For temperatures in between, where we have the chance of both stability and changeability, we observe more interesting dynamics' (1986:128). The 'more interesting dynamics' result in structures of sufficient complexity and variety that they might mirror the generative complexity of a relatively simple cognitive system.

Computational temperature

Langton's simulations suggest a model in which the cognitive system works most effectively over a narrow range of computational temperatures. Low temperatures lead to relatively low levels of inactivity and simple patterns. Higher levels of activity and more complex patterns are obtained by increasing computational temperatures. The variations in the metastable regimes between the stable quiescent state (at low temperatures) and the chaotic state (at high temperatures) generate the most interesting and creative activity. Increases in computational temperature lead to increasingly complex patterns until the system breaks down by transitioning into chaotic activity. It is in this metastable region prior to chaos that the mapping from micro states to macro behavior remains coherent yet is sufficiently complex to pose a challenge to understanding the nature of that mapping.

Learning and memory

In our initial presentation of the required features for associative memory we noted that the system must be able to generate ordered states from homogeneity but that it must also be able to select those emergent states more readily on subsequent occasions (that is it must learn). One challenge that remains for the cellular automata model is to show that the system can learn. Possibly as a function of repeated exposure to inputs or to repeated activation of the system the transition function might evolve from one that produces no ordered states to one that produces useful patterns. Langton's system is, however, a model of a closed system and, from the perspective of our earlier discussion, flows of matter as well as of energy (i.e., an open system) must be modeled to permit the emergence of a memory.

Order from disorder in neural networks: Cognition in an open system?

An associative memory model proposed by Baird (1986) has distinctive spatial patterns of neural activity representing modal states for distinctive recognition and response behaviors. The empirical data that underlie this model are EEG recordings collected by Freeman and his associates from the olfactory bulb of the rabbit (see Freeman and Skarda 1985; Skarda and Freeman 1987). The problem faced in this work was to isolate distinctive patterns of behavior that could serve to identify invariant odor classes.

Associative memory in the olfactory bulb

The EEG recordings showed that odor recognition is accompanied by space-time patterns of peak RMS neural activity distributed throughout the bulb. Skarda and Freeman (1987) concluded that every neuron in the olfactory bulb participated in every discrimination, and that discrimination between odors was based on the assembly of different patterns of neural intensity during inhalation. Baird (1986) took this as evidence that the olfactory bulb has no feature processing subsystems and can best be described as an associative memory in which the emergent spatial inhomogeneity results from the pattern of strengths in synaptic connections.

The fact that different space-time patterns of peak RMS activity are generated in response to different odors is suggestive of behavior like that modeled with the interactive activation and competitive networks of McClelland and Rumelhart (1988). In those networks a single pattern of connections will settle into different states based on the pattern of input strengths. The emergence of different patterns depends considerably on the interaction between excitatory and inhibitory processes (fig. 2). By virtue of this interaction a number of different latent states can coexist within the same set of connections. In the olfactory bulb these constitute the set of latent activity patterns that represent the odor memories or the invariant odor classes to be recognized.

From observations on the olfactory bulb Baird (1986) developed a view of the entire cortex as a set of associative memories interconnected in parallel. Some of the challenges for this view are to specify how a specific pattern of activity is selected, how the system transitions between states, how activity permeates through the system to select particular sets of patterns in the various associative memories that correspond to particular behaviors, how to account for both the rapid changes between some cognitive states and the relatively long-term persistence of others, and how new patterns (including the first) are established.

Self-organization in the olfactory bulb: Order through heating

The recognition of a learned odor is postulated to occur via a process of order emerging from disorder, or pattern from homogeneity as in the Rayleigh-Benard convection. EEG data show a brief transition from a low-level state of irregular activity to the high amplitude pattern associated with a specific odor. Thus, recognition occurs in an associative memory when the system is driven from a homogeneous state of low-level, stochastic neural activity beyond the threshold of stability, through a bifurcation which places it in the basin of an attractor. Within the olfactory bulb the order parameter that produces the symmetry break via the bifurcation is thought to be the higher energetic state that accompanies inhalation. More generally the order parameter may be associated with energy dissipation through the system; that is the metabolization of glucose transported by blood flows (Iverson 1979: 70).

Macro selection of odor basins via biasing inputs

Given the existence of multiple latent states in close competition some mechanism is required for reliable selection of an appropriate modal state. In a system such as the Rayleigh-Benard convection the order parameter drives the system to the point of bifurcation where a stochastic fluctuation is amplified to lead the system into a new modal state. Within an associative memory stochastic fluctuations presumably have some effect but for reliable classification of inputs those inputs must act as low energy biases to lead the system into the required attractor basin as it is forced from its state of stable equilibrium. In that sense the model employs a mechanism that permits an adaptive versus a spontaneous symmetry break from the homogeneous state.

Transitions between odor basins

A reset mechanism is required to allow the system to settle into different modal states that constitute recognition of different inputs. Baird (1986) suggests relaxation back to the resting state via inverse bifurcation. While he discounts the possibility of a catastrophic transition induced by an input fluctuation as a means of placing the system in a different attractor basin such catastrophic transitions are found in many physical systems. This would, however, require a higher level of input energy than is required for selection of a new state via the inverse bifurcation route.

Whether input changes can produce energy fluctuations needed for catastrophic transitions is unknown but such a process would be consistent with the problem solving experience of transitioning suddenly from an incorrect to a correct solution while continuing to maintain a high level of effort. Transitions via inverse bifurcation would be more like the process of achieving a correct solution by putting the problem to one side for a time. Inverse bifurcation and catastrophes are not mutually incompatible processes and both may have a role to play.

Categorical odor perception

The appeal to attractors and basins of attraction allows inputs to be characterized as noisy or incomplete versions of an attractor. Different inputs that drive the system into the same attractor basin (and the number of these is potentially infinite) and therefore cause the system to collapse onto the same attractor will be classified as identical. This view of how the macro response of an associative network is assembled is consistent with recent discussions of categorical perception (e.g., Harnad 1987). In addition, Skarda and Freeman (1987) have observed the occasional failure to recognize odors. This appears to result in a disorderly or chaotic attractor that cannot be classified, and may generally be associated with indecision (Baird 1986).

Emergence of distributed, high-order patterns

Given a system of multiple, interconnected associative networks, the output from one system that constitutes an input for a second system can influence the state of the second system by pushing it into a different attractor basin. It is by this process that change can permeate through the system and that different combinations of associative states can emerge. As each and every associative network within a mature adult will have multiple latent states the potential combinatorial complexity is enormous. A reasonable extension of Baird's views would have some associative networks (e.g., those most directly connected to the environment) transitioning between states at a relatively high rate (2 to 3 times per second). Others (e.g., those that influence persistent goal-directed activity) are likely to transition between states much less frequently, possibly at intervals of minutes, hours, days, and years.

Learning and differentiation

The differentiation of attractors is accomplished via the modification of the coupling within the network. As is consistent with learning rules employed in PDP, excitatory connections may be strengthened by concurrent activity. The particular pattern established for a new input is most likely an arbitrary function of the particular peaks in stochastic activity at the time of input and of the activity generated by the input. Where specific inputs are repeated with sufficient frequency, patterns will emerge once connection strengths exceed a critical threshold much in the manner that an arbitrary pattern emerges in the termite nest-construction field. Once connection strengths have been established the input activity must emerge as a more powerful lead in the direction taken by the subsequent activity within the network.

Nevertheless, the EEG data of Skarda and Freeman (1987) indicate that patterns are not fixed. On return to a previously learned odor (apparently some weeks later) the pattern could differ from the one previously evoked by that odor. In general, the learning of a new odor appeared to result in dynamic reconstruction of the whole patterned response set. It is this process of dynamic reconstruction that may permit progressive differentiation of stimulus information.

Associative memory as a dissipative structure

Baird's associative memory model may be viewed as connectionist but as one that avoids some of the major difficulties we see in the general trend of PDP modeling. It addresses the problems of mappings between micro and macro states in a manner consistent with the principles of self-organization. Baird's view of cognitive process as based on dissipative dynamics provides a perspective that we view as critical to a model of cognition that heeds the realities of biological processes.

Learning and memory

The central nervous system is open to flows of both energy and matter. In that sense it is more like the termite nest construction system than like any of the other thermodynamic systems we have discussed in this paper. Not only is it possible to select modal states but experience can create a preferred status for particular states. That is, a state that has been achieved once can be achieved more easily in the future. This is a general characteristic of open systems; residues accumulate as evidence of past dynamic activity.

Associative memory revisited

A viable associative memory system must have mechanisms that permit macro states to be selected and assembled on a first occasion and that allow some fraction of the potential states to emerge with a preferred status as a result of learning. It must have a mechanism for selection that may be part stochastic, but should be responsive to external input, and it must have a mechanism for cyclically assembling, disassembling, and reassembling macro states. We find that closed systems (Hopfield and Tank 1986; Langton 1986; Smolensky 1987), and open systems (Baird 1986; Kugler and Turvey 1987) all exhibit some of the necessary features but an open systems theory as promoted by Bertalanffy (1975) and further developed by Prigogine and his colleagues (e.g., Prigogine and Stengers 1984) shows most promise for simultaneously satisfying all requirements.

The use of termite nest construction as a paradigmatic example of self-organization in an open system has value in that it is possible to identify mechanisms that lead to macroscopic order. Our discussion of termite nest construction allowed us to isolate the roles of stochastic fluctuation, perceptual sensitivity, and autocatalytic symmetry breaking as they act through successive orders of organization. This example also illustrates the nature of memory as a physical instantiation of a symbolic, rate-independent process with its origins in a dynamic, rate-dependent process.

Baird's (1986) account shows that multiple preferred states can be established and can continue to coexist within the same distributed network. Given the relatively meager knowledge about the functioning of neural networks it is not possible to provide a clear account of the mechanisms involved in the emergence of preferred states from the multiple potentialities, in their assembly, disassembly and reassembly, or in their selection. A primary claim advanced here is that whatever the mechanisms they will be consistent with an open systems view.

Our view is in clear contrast to one offered by Johnson-Laird (1983: 399–406) who views self-organization as characteristic of noncognitive or nonanticipatory systems. He argues that cognitive activity relies on a high-level, representational model of the world; one that is based in an arbitrary symbolic notation. Nevertheless, self-organization is not antithetical to anticipation or intention (Shaw and Kinsella-Shaw 1988) and one primary claim for this paper is that a self-organizational account of cognition deserves serious consideration.

Rules versus laws

Organization may be achieved via the implementation of rules that specify how degrees of freedom are to be constrained. The serial, digital computer accomplishes this by representing the desired output in the program code. Connectionists argue that their systems self-organize which implies that the features of organization in the macro behavior are not represented in the program code. Nevertheless a connectionist system has to be appropriately tuned for it to generate interesting behavior. For example, the inhibitory and excitatory processes of an Interactive Activation–Competition network (McClelland and Rumelhart 1988) must be set appropriately for the system to model even simple forms of human perceptual or judgmental behavior. For such a network to provide a compelling explanation of human behavior some principled basis for selecting the architecture and its parameters (i.e., the symbolic constraints) must be established.

The construction of termite nests is influenced by dynamical laws. It is presumably possible to develop a rule-based simulation of termite nest construction in which the appropriate macro properties were specified but what such a simulation would offer in terms of understanding or explanation is not clear (fig. 9). Similarly, a rule-based

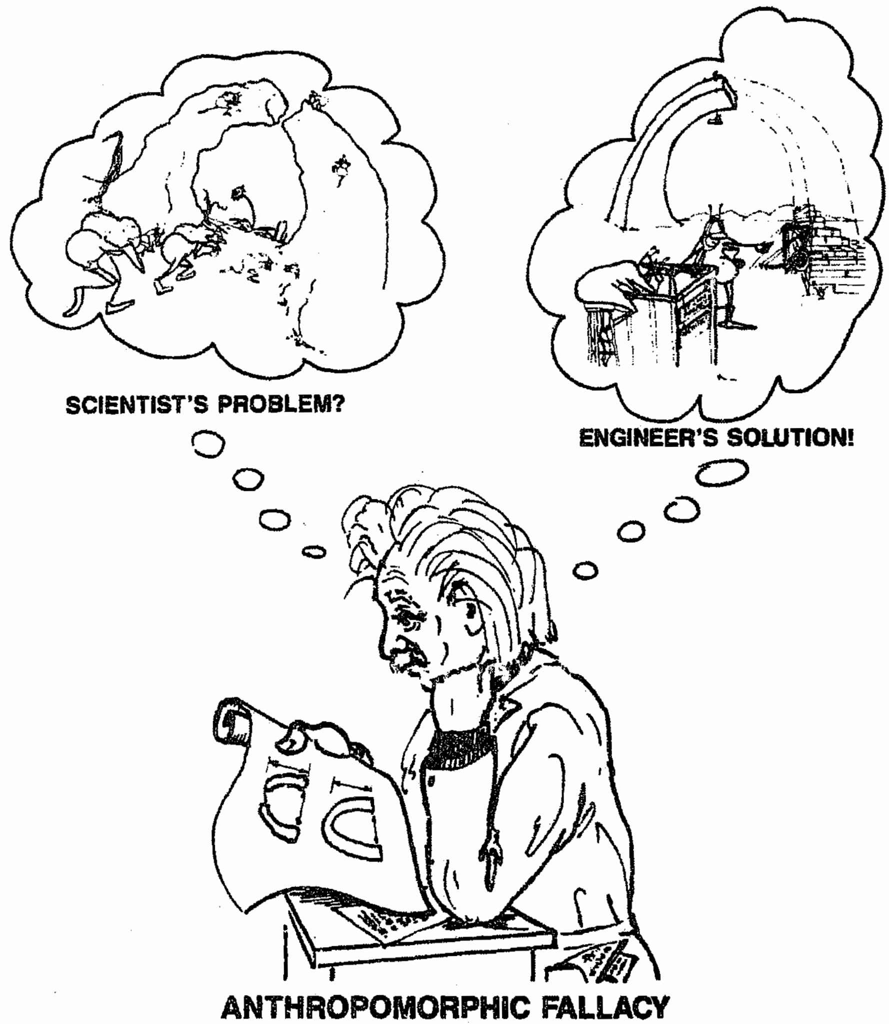

Fig. 9. The modeling problem. A rule-based simulation of an organized behavior does not clarify the nature of the mechanisms that underlie the organization.

simulation of cognition is unlikely to be enlightening. We need an approach to modeling that honors lawful dynamical processes. No rule-based system that establishes control through rate-independent symbolic processes can provide a compelling model of human cognitive behavior without a principled account of the origin of those rules. Additionally, there needs to be some account of how the appropriate rules could produce complex patterns of organization the details of which are not specifically represented in the code. We propose that the most likely source is in the dynamic rate-dependent processes that are given expression in the formulation of natural laws.

Law versus analogy

It is acknowledged, within the connectionist literature, that thermodynamic law is exploited for the analogies [3] it offers the study of cognition (Hopfield and Tank 1986; Norman 1987; Smolensky 1987). Indeed, the use of a freezing principle to describe the process of achieving a stable cognitive state must be viewed as an analogy because actual freezing of a biological system would lead to death. However, our reference to termite nests should not be taken as analogical. This is a prototypical example of self-organization in a biological system. The value of referring to termite nest construction when the concern is specifically with cognition is that it is possible to explore interesting facets of biological self-organization and to generate interesting hypotheses prior to the more challenging task of exploring self-organization in neural systems.

It would be a mistake to view the application of thermodynamic law by Haken (1981), Atkins (1984), or Prigogine and Stengers (1984) to an explanation of self-organization in physical or behavioral systems as an analogy. In their terms, thermodynamic law is a description of universal and natural regularities that have causal potency. It is in the Second Law that lies the seeds of change. The Second Law is behind the examples of self-organization that have been outlined in this paper and, for some, it accounts for the creation and decay of all structure in the universe (Atkins 1984; Haken 1981; Prigogine and Stengers 1984). The Rayleigh-Benard convection is an ordered structure that is created by a dissipative dynamic. Similarly, the pattern of organization in the termite pheromone field is sustained by a flow of energy and matter as is consistent with the Second Law (Deneubourge 1977). These systems may be contrasted to a digital computer which is a dualistic machine; energy flows serve to sustain the machine in a state that maintains its capability for work but the energy flows have no influence on the nature of the symbolic processing.

The patterns that emerge from the influence of the Second Law are referred to as dissipative structures (Prigogine and Stengers 1984),

[3] Most frequently referred to as metaphor in the PDP literature. A metaphor is an elegant and creative expression of an idea; a figure of speech (e.g., the cat is out of the bag), while an analogy expresses similarity on some dimension. Analogy is the correct characterization of thermodynamic principles as employed in PDP modeling for the simulation of cognition.

which in effect are local abatements of entropy (disorder, homogeneity). Dissipative structures are Nature's sleight of hand; they are local structures in which order is created in a manner that speeds the global progress towards universal entropy. Kugler and Turvey (1987), for whom cognition is also a dissipative structure, employ the Second Law of thermodynamics as a law in contrast to an analogy, and are in concert with Atkins (1984), Haken (1981), and Prigogine and Stengers (1984) in viewing it as universally responsible for order. From that perspective the Second Law has certain imperatives that cannot be ignored. Any account of cognition that is not consistent with the Second Law is necessarily flawed.

On the other hand, the use of thermodynamic law as a source of analogies carries no such imperatives. The value of an analogy lies in its heuristic potential and in the ability of the user to exploit that potential. Analogies offer considerable freedom. They may be exploited in any way they are found to be useful and much of the value of a specific analogy may stem from the creativity of the user rather than from the intrinsic power of the analogy. Nevertheless, the freedom offered by the use of analogies can also be viewed as a lack of constraint and this may account for the seemingly arbitrary proliferation of PDP models. Cognitive science has relied heavily on analogies in the past; a strategy that may have contributed to a situation in which it is vulnerable to the charge that 'each new experimental finding seems to require a new theory' (Norman 1987: 535). It is nevertheless ironic that this charge emerges from the connectionist literature because the arbitrary proliferation of models is one glaring problem with the PDP enterprise.

The claim offered here is that analogy is a poor basis for a paradigm shift, or even for a scientific thrust (also see Bertalanffy 1968: 84–85). There is no doubt that analogy can be useful for generating or communicating ideas, understanding, or hypotheses, but the development of any scientific endeavor must ultimately be based on laws. To relegate laws to a subsidiary role will lead to unprincipled distinctions and theoretical elaborations that contribute little to the progressive construction of a useful body of knowledge.

Summary

Connectionists, by their concern with self-organization, micro and macro states, and the mapping between them have introduced a valuable emphasis to the study of human behavior. Consideration of the mapping between micro and macro states is essential to the understanding of complex, nonlinear systems, of which humans, collectively or individually, are prime examples. An appeal to self-organization offers an alternative to classical reductionism and, when understood in detail, a solution to the homunculus problem. Nevertheless, the view of self-organization presented in the PDP literature is limited, particularly in the nature of the physical examples that guide the modeling efforts. While computational temperature would appear to provide an effective order parameter for the emergence of structure in cognition, the reliance on a heating principle for the emergence of diverse and creative patterns of behavior will be more productive than reliance on a cooling principle. The concept of a dissipative structure within the framework of an open systems account has much to offer the study of associative memory or of cognition.

In addition, analogy is a precarious basis for a paradigm shift, or even for the less ambitious tasks of theorizing and model development. The use of thermodynamic law as a source of analogies denies the imperatives that would flow from giving it full force as a lawful basis for the emergent structure of cognition. A lawful account of cognition based on thermodynamic law is likely to avoid the unprincipled proliferation of explanations and models and is likely to lead to a more coherent development of our science. The challenge remains to turn these ideas into an account of the human behaviors that are central to adaptive functioning in a complex world. This is a challenge that the dominant forces in 100 years of experimental and cognitive psychology have so far failed to tackle in any substantive way.

References

Adams, J.A., 1971. A closed loop theory of motor learning. Journal of Motor Behavior 3, 111–149.

Atkins, P.W., 1984. The second law. New York: W.H. Freeman.

Baird, B., 1986. Nonlinear dynamics of pattern formation and pattern recognition in the rabbit olfactory bulb. Physica 22D, 150–175.

Ballard, D.H., 1986. Cortical connections and parallel processing: Structure and function. The Behavioral and Brain Sciences 9, 67–120.

Berge, P., Y. Pomeau and C. Vidal, 1984. Order within chaos: Towards a deterministic approach to turbulence. New York: Wiley.

Bertalanffy, L. von, 1968. General system theory. New York: George Braziller.

Bertalanffy, L. von, 1975. Perspectives on general systems theory: Scientific philosophical studies. New York: George Braziller.

Casti, J.L., 1989. Paradigms lost: Images of man in the mirror of science. New York: William Morrow.

Crick, F.H.C., 1979. 'Thinking about the brain'. In: Scientific American. The brain. New York: W.H. Freeman. pp. 130–137. (Reprinted from Scientific American, 1979.)

Crick, F. and C. Asanuma, 1987. 'Certain aspects of the anatomy and physiology of the cerebral cortex'. In: J.L. McClelland and D.E. Rumelhart (eds.), Parallel distributed processing. Explorations in the microstructure of cognition. Vol. 2: Psychological and biological models. Cambridge, MA: The MIT Press.

Deneubourge, J.L., 1977. Application de l'ordre par fluctuation a la description de certaines étapes de la construction du nid chez les termites. Insectes Sociaux, Journal International pour l'Etude des Arthropodes Sociaux 24, 117–130.

Freeman, W.J. and C.A. Skarda, 1985. Spatial EEG patterns, non-linear dynamics and perception: The neo-Sherringtonian view. Brain Research Reviews 10, 147–175.

Gibson, J.J., 1979. The ecological approach to visual perception. Boston, MA: Houghton Mifflin.

Grasse, P.P., 1959. La reconstruction du nid et les coordination interindividuelles chez Bellicositermes natalensis et cubitermes sp. La théorie de la stigmergie: essai d'interprétation des termites constructeurs. Insectes Sociaux, Journal International pour l'Etude des Arthropodes Sociaux 6, 127.

Haken, H., 1981. The science of structure: Synergetics. New York: Van Nostrand Reinhold.

Harnad, S., 1987. 'Category induction and representation'. In: S. Harnad (ed.), Categorical perception. Cambridge, MA: Cambridge University Press.

Hinton, G.E. and T.J. Sejnowski, 1987. 'Learning and relearning in Boltzmann machines'. In: D.E. Rumelhart and J.L. McClelland (eds.), Parallel distributed processing. Explorations in the microstructure of cognition. Vol. 1: Foundations. Cambridge, MA: The MIT Press.

Hopfield, J.J. and D.W. Tank, 1986. Computing with neural circuits: A model. Science 233, 625–633.

Iverson, L.L., 1979. 'The chemistry of the brain'. In: Scientific American. The brain. New York: W.H. Freeman. pp. 70–81. (Reprinted from Scientific American, 1979.)

Johnson-Laird, P.N., 1983. Mental models. Cambridge, MA: Harvard University Press.

Kugler, P.N., 1986. 'A morphological perspective on the origin and evolution of movement patterns'. In: M. Wade and H.T.A. Whiting (eds.), Motor development in children: Aspects of coordination and control. Boston, MA: Martinus Nijhoff.

Kugler, P.N. and M.T. Turvey, 1987. Information, natural law, and the self-assembly of rhythmic movement. Hillsdale, NJ: Erlbaum.

Langton, C.G., 1986. Studying artificial life with cellular automata. Physica 22D, 120–149.

May, R.M., 1976. Simple mathematical models with very complicated dynamics. Nature, 459–467.

McClelland, J.L. and D.E. Rumelhart, 1987. Parallel distributed processing. Explorations in the microstructure of cognition. Vol. 2: Psychological and biological models. Cambridge, MA: The MIT Press.

McClelland, J.L. and D.E. Rumelhart, 1988. Explorations in parallel distributed processing. Cambridge, MA: The MIT Press.

McClelland, J.L., D.E. Rumelhart and G.E. Hinton, 1987. 'The appeal of parallel distributed processing'. In: D.E. Rumelhart and J.L. McClelland (eds.), Parallel distributed processing. Explorations in the microstructure of cognition. Vol. 1: Foundations. Cambridge, MA: The MIT Press.

Metropolis, N., A. Rosenbluth, M. Rosenbluth, A. Teller and E. Teller, 1953. Equations of state calculations for fast computing machines. Journal of Chemical Physics 21, 1087–1091.

Norman, D.A., 1987. 'Reflections on cognition and parallel distributed processing'. In: J.L. McClelland and D.E. Rumelhart (eds.), Parallel distributed processing. Explorations in the microstructure of cognition. Vol. 2: Psychological and biological models. Cambridge, MA: The MIT Press.

Pattee, H.H., 1979. The complementarity principle and the origin of macromolecular information. BioSystems 11, 217–226.

Pew, R.W., 1970. Toward a process oriented theory of human skilled performance. Journal of Motor Behavior 11, 8–24.

Prigogine, I. and I. Stengers, 1984. Order out of chaos. New York: Bantam Books.

Rosenblatt, F., 1959. 'Two theorems of statistical separability in the perceptron'. In: Mechanisation of thought processes: Proceedings of a symposium held at the National Physical Laboratory, November 1958, Volume 1. London: HM Stationery Office. pp. 421–456.

Rosenblatt, F., 1962. Principles of neurodynamics. New York: Spartan.

Rumelhart, D.E. and J.L. McClelland, 1987. Parallel distributed processing. Explorations in the microstructure of cognition. Vol. 1: Foundations. Cambridge, MA: The MIT Press.

Schmidt, R.A., 1975. A schema theory of discrete motor skill learning. Psychological Review 82, 225–260.

Schneider, W., 1987. Connectionism: Is it a paradigm shift for psychology? Behavior Research Methods, Instruments, and Computers 19(2), 73–83.

Shaw, R. and J. Kinsella-Shaw, 1988. Ecological mechanics: A physical geometry for intentional constraints. Human Movement Science 7, 155–200.

Skarda, C.A. and W.J. Freeman, 1987. How brains make chaos in order to make sense of the world. Behavioral and Brain Sciences 10, 161–195.

Smolensky, P., 1987. 'Information processing in dynamical systems: Foundations of harmony theory'. In: D.E. Rumelhart and J.L. McClelland (eds.), Parallel distributed processing. Explorations in the microstructure of cognition. Vol. 1: Foundations. Cambridge, MA: The MIT Press.

Smolensky, P., 1988. On the proper treatment of connectionism. Behavioral and Brain Sciences 11, 1–74.

Wolfram, S., 1984. Universality and complexity in cellular automata. Physica 10D, 1–35.