BEHAVIORAL AND BRAIN SCIENCES (1987) 10, 407–465
Printed in the United States of America

Levels of modeling of mechanisms of visually guided behavior

Michael A. Arbib
Departments of Computer Science, Neurobiology, Physiology, and Biomedical Engineering, University of Southern California, Los Angeles, Calif. 90089

Abstract: Intermediate constructs are required as bridges between complex behaviors and realistic models of neural circuitry. For cognitive scientists in general, schemas are the appropriate functional units; brain theorists can work with neural layers as units intermediate between structures subserving schemas and small neural circuits.

After an account of different levels of analysis, we describe visuomotor coordination in terms of perceptual schemas and motor schemas. The interest of schemas to cognitive science in general is illustrated with the example of perceptual schemas in high-level vision and motor schemas in the control of dextrous hands.

Rana computatrix, the computational frog, is introduced to show how one constructs an evolving set of model families to mediate flexible cooperation between theory and experiment. *Rana computatrix* may be able to do for the study of the organizational principles of neural circuitry what *Aplysia* has done for the study of subcellular mechanisms of learning. Approach, avoidance, and detour behavior in frogs and toads are analyzed in terms of interacting schemas. Facilitation and prey recognition are implemented as tectal–pretectal interactions, with the tectum modeled by an array of tectal columns. We show how layered neural computation enters into models of stereopsis and how depth schemas may involve the interaction of accommodation and binocular cues in anurans.

Keywords: cognition; computation, cooperative; connectionism; frogs; motor patterns; neural modeling; neuroethology; prey recognition; schemas; stereopsis; tectum

1. Multiple levels of analysis

Recently, the domination of cognitive science by the serial computational paradigm has been broken by research on "connectionism" and "parallel distributed processing" (see Rumelhart and McClelland [1986] for a collection of papers; and Arbib [1987] for a historical perspective linking this field to earlier work in brain theory and adaptive networks). The stress, however, has been on models with only two levels: overall behavior and "neuronlike" elements. In this target article, I argue for the utility of *schemas* as units of computational analysis intermediate between behavior and neuron, while stressing that brain theory is not just connectionism: Incorporating data from neuroanatomy and neurophysiology requires considerations that go well beyond the development of cognitive models "in the style of the brain."

In this target article, I will discuss levels of analysis that can bridge between neurons and behavior. Schemas and their use in models of vision and visually guided behavior will be introduced to suggest their utility for cognitive science in general. *Rana computatrix,* a family of computational models of neural circuitry underlying visually guided behavior in frog and toad, will be used to illustrate strategies for moving from schema-theoretic models to testable models of neural circuitry.

I argue that top-down analysis of animal behavior should yield a functional model of how that behavior is achieved through the interaction of simultaneous computing agents called *schemas.* These must be "large" enough to allow an analysis of behavior at or near the psychological level. In brain theory, we must also turn to lesion studies (at a gross level) and neurophysiology and neuroanatomy (at a finer level) to suggest how the schemas might be implemented in the brain of the organism under study.

Figure 1 emphasizes that our choice of levels of analysis can be functional or structural. A top-down analysis starts with picking some overall animal behavior for study – but even this selection is theory laden, and what seems like a natural choice for behavior may prove on further study not to be unitary at all. For this reason, I regard Marr's (e.g., 1982) notion of an independent computational level of analysis as mistaken – for example, one cannot give an a priori analysis of depth perception because different animals (or different subsystems of a given animal) may make different use of visual cues that cannot be discovered until "implementational details" (the data of neuroscience!) are taken into account. We will suggest that depth perception in the frog and toad is not unitary; quite different mechanisms are involved in depth perception for prey and for barriers. If Marr is interpreted as likewise allowing such an account, then his levels seem to add nothing to earlier elucidation of the virtues of a functional analysis (Arbib 1975). Proceeding with such an analysis, we explain behaviors in terms of interactions among functional units (called "schemas") whose nature will be discussed in more detail below. If we start with a structural analysis, the brain region defined by cytoarchitectonics or input/output pathways provides the large-

0140-525X/87 $5.00+.00

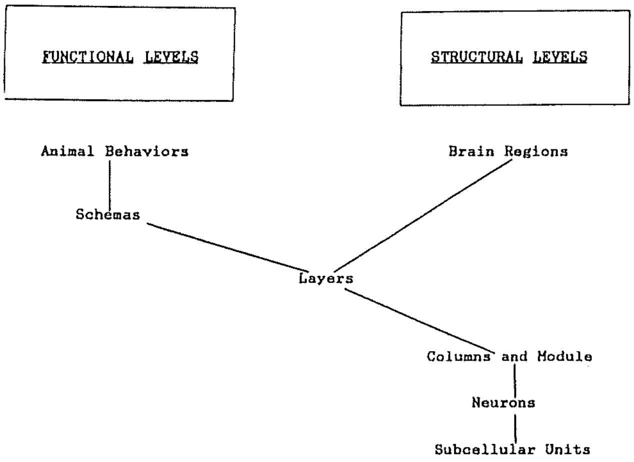

Figure 1. Levels of functional and structural analysis.

scale framework. Then, as we try to make the bridge from large unit to neuron, we can elaborate our models in terms of such intermediate constructs as arrays (layers), columns, or modules, which can be characterized either functionally or structurally. The eventual goal, of course, is for functional and structural analyses to be rendered congruent.

For many behaviors, analysis at the level of single neurons may be superfluous. However, even when this further level is appropriate, there remains the question of how detailed the analysis of each neuron must be. Consider Figure 2a, which shows Székely and Lázár's (1976) Golgi study of the frog tectum. There is a diversity of neurons and a richness of structure, yet the diagram tells us little about the detailed connectivity needed to relate structure to function. One can also have too much detail, however, as can be seen in the reconstruction from electron microscopy in Figure 2b. Detailed studies of single neurons are undeniably important, but these should be used to fine-tune more economical "black-box" models of neurons. It is these simplified model neurons that serve as the units in our models of large networks. The models described in Section 4 offer one approach to such a simplification. As detailed in Appendix C, we have described the behavior of each neuron by a simple difference or differential equation (linear in terms of the synaptically weighted input values), representing the firing rate of each cell by a simple nonlinear transformation of the membrane potential.

In modeling the circuitry of the tectum (Figure 2d, based on Lara et al. 1982), we were guided by Székely and Lázár's "tectal column," Figure 2c. Unlike in analyses of mammalian cortex, this terminology does not imply that adjacent columns must have different response properties. Rather, the tectal column is simply a basic unit of vertical organization – a "tile" that can be repeated again and again to approximate the tectal neural net. In Section 4, we will outline how an array of such tectal columns has been used in modeling prey-predator discrimination.

Quite apart from the details of its specification, a neuron may serve as a functional unit, a structural unit, or both. We can determine what are units in the brain physiologically (e.g., by electrical recording) or anatom-

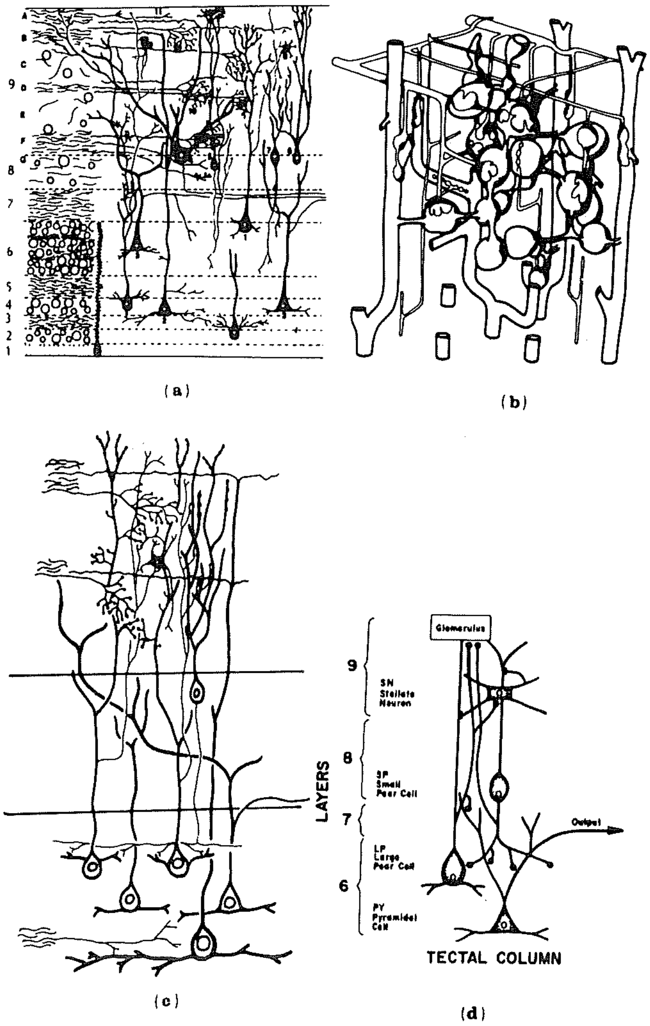

Figure 2. (a) Diagrammatic representation of the lamination and the representative types of neurons of the optic tectum. Numbers on the left indicate the different tectal layers. Numbered cell types are as follows: (1) large pear-shaped neuron with dendritic appendages and ascending axon; (2) large pear-shaped neuron with dendritic collaterals; (3) large pyramidal neuron with efferent axon; (4) large tectal ganglionic neuron with efferent axon; (5–6) small pear-shaped neurons with descending and ascending axons respectively; (7) bipolar neuron; (8) stellate neuron; (9) amacrine cell; (10) optic terminals; (11) assumed evidence of diencephalic fibers (from Székely & Lázár 1976). (b)Details of synaptic interaction of dendritic appendages, which exceed current models in intricacy (from Székely & Lázár 1976). (c) Schematic representation of a tectal column (from Székely & Lázár 1976). (d) Neurons and synaptology of the model of the tectal column. The numbers at the left indicate the different tectal layers. The glomerulus is constituted by the LP and SP dendrites and recurrent axons as well as by optic and diencephalic terminals. The LP excites the PY, the SN, and the GL, and is inhibited by the SN. The SP excites the LP and PY cells, and it sends recurrent axons to the glomerulus; it is inhibited by the SN. The SN is excited by LP neurons and diencephalic fibers, and it inhibits the LP and SP cells. The PY is activated by the LP, SP, and optic fibers, and is the efferent neuron of the tectum (Arbib 1982b).

ically (e.g., by staining). In many regions (but not all) we know which anatomical entity yields which physiological response. For example, the identification of the physiological responses recorded from the tectum during the facilitation of prey-catching behavior shown in Figure 13d with specific anatomically defined cells in Figure 2a is a current research matter (as initiated by Antal et al. 1986; Matsumoto et al. 1986) that is not yet reflected in modeling studies. Nonetheless, hypotheses had to be made in formulating and testing our models. For the model in Figure 2d, other choices also had to be made. Figure 2a shows two types of output cells in the tectum: the pyramidal cell and the large tectal ganglionic neuron. The model in Section 4 assumes only that the output of the former is relevant to the phenomenon of prey selection. Our models must be adaptable to phenomena that depend on ganglionic output.

It is for such reasons that our modeling methodology explores a variety of different connectivities in some overall paradigm of brain function. We set up a *model family* that allows us to experiment with a number of different connectivities and parameter settings for the cells of the model. This leads to a style of *incremental modeling*. For example, the first model of a "tectal column" was introduced to explain certain facilitative effects in prey-catching behavior; a linear array of such columns was then used to model certain data on size dependence of prey-catching activity in toads; inhibition from the pretectum to such an array was then introduced to model the behavior of an animal confronted with more than one prey stimulus. Elsewhere (Arbib 1982a) I have referred to incremental modeling as "evolutionary modeling" because these models form three stages in an evolutionary sequence (although it is the modeler rather than nature who introduces the variations and does the selection) for what I have dubbed *Rana computatrix*, "the computational frog" – our developing model of the neural circuitry underlying visuomotor coordination in the frog and toad.

2. Perceptual and motor schemas

Although there are formal characterizations of programs in FORTRAN, or LISP, or PASCAL, there is no single definition that encompasses all programs: serial, parallel, or concurrent, recursive or nonrecursive, and object-based or not. Despite this, computer scientists recognize commonalities that allow new program concepts to be built on old ones. In the same spirit, our work on schemas has not yet yielded a single formalism, but contributes to the evolution of a theory of schemas. The schemas are programs developed to satisfy the following criteria (Arbib 1981):

(a) Schemas represent perceptual structures and distributed motor control. Hypotheses about how behavior is generated are formulated in terms of assemblages of perceptual schemas and coordinated control programs of motor schemas.

(b) Schemas can be instantiated. Given a schema that represents generic knowledge about some domain of interaction (e.g., a chair and how to sit on it), we need several schema instantiations, each suitably tuned, to subserve our perception of several instances of that domain.

(c) Like procedures or programs, schemas can be combined to form new schemas. In particular, a given schema may be instantiated many times within a larger schema instantiation. In serial programs, no matter how many the subroutine calls in a program, only one call will be executed at any given time. By contrast, schema instantiations are inherently concurrent. We postulate that, unlike serial computers, the brain can support the concurrent activity of many schemas for the recognition of different objects and for the planning and control of different activities.

(d) As constrained by criteria (a)–(c) above, schemas can enter into theories of AI, perceptual robotics, or cognitive psychology. To qualify as part of brain theory (computational neuroscience), schemas must be further linked to a theory about neural localization or, at an even finer-grained level, to their implementation in neural circuitry.

The schema notion will be made explicit in three applications: frog visuomotor coordination (Section 3), machine vision (Appendix A), and robotics (Appendix B). In Arbib (1975) I introduced the terms "perceptual schema" and "motor schema." At about the same time, Minsky (1975) established the term "frame" in the AI community. I did not abandon the term "schema," however, for three reasons: (a) There was a continuity between my concerns and those of Bartlett (1932), Head and Holmes (1911), Piaget (1971), Schmidt (1975), and other writers who had already used the term "schema." (b) The term frame already had two uses – in "the frame problem" of AI (when an operator is applied, how does one keep track of the changes not explicitly specified in the description of the operator?), and in the "frame" as the unit picture in a movie that is a candidate for analysis by a (machine) vision system. (c) Frames (Minsky 1975) and scripts (Schank & Abelson 1977) were developed with an emphasis on representing the overall framework in which a situation is to be analyzed; my work, on the other hand, emphasizes the constructive aspect, with perceptual and motor schemas being combined to form schema assemblages and coordinated control programs that can themselves form new schemas to be combined into yet larger schemas. Despite these differences, I see developments in frame theory and script theory as part of the same enterprise as schema theory.

The internal model of the world (Craik 1943; Gregory 1969; MacKay 1966; Minsky 1961) must be built of units, each of which corresponds roughly to a *domain of interaction*, which may be an object in the usual sense, an attention-riveting detail of an object, or some domain of social interaction (Arbib 1975). The system determines whether a given domain of interaction is present in the environment using a *perceptual schema*. The state of activation of the schema then determines the strength of the hypothesis that what the schema represents is indeed present. Other schema parameters represent properties such as size, location, and motion of the perceived object. A *schema assemblage* – an assemblage of *instantiated* perceptual schemas – provides a short-term memory (STM) combining an estimate of environmental state with a representation of goals and needs. Long-term memory (LTM) is provided by the stock of schemas from which STM can be assembled. (Appendix A describes a computer system for constructing a schema assemblage to represent a photograph of an outdoor scene [Weymouth 1986].) New sensory input and internal processes update the schema assemblage. This can itself be action-dependent, as in actively palpating an object. The internal state is also updated by knowledge of the state of execution of current *plans*. We hypothesize that these plans are made up of *motor schemas*, which are akin to control systems but can be combined to form coordinated control programs to control the phasing of various patterns of coactivation.

Biological control theory usually studies neural circuitry specialized for the control of a specific function, be it the stretch reflex (Stein 1982) or the vestibulo-ocular reflex (Robinson 1981). Yet most behavior involves complex sequences of the coordinated activity of a number of control systems. I have introduced the notion of *coordinated control program* (Arbib 1981) as a combination of control theory and the computer scientist's notion of a program suited to the analysis of the control of movement. Such a program can control the time-varying interaction of a number of schemas. In the diagrams depicting such a program, there are lines representing both transfer of activation and transfer of data.

Although we will emphasize visuomotor coordination in the frog and toad in later sections, our first example of a coordinated control program is based on the study of human reaching to grasp an object. The top half of Figure 3 (Jeannerod 1981) shows a collage of hand shapes as the hand moves from its initial position to pick up a ball. As the hand moves, it is *preshaped* so that when it has almost reached the ball, it has the right shape and orientation to enclose the ball prior to gripping it firmly. The lower half of the figure indicates with a dot the position of the tip of the thumbnail in consecutive frames of a movie. Examining the spacing of these dots, we can see that the movement can be broken into two parts: a fast initial movement, and a slow approach movement. Moreover, Jeannerod has shown that the transition from the fast to the slow phase of the hand transfer movement is coupled with a transition from the preshape of the hand itself to the closing in of the fingers so that touch can take over in controlling the final grasp.

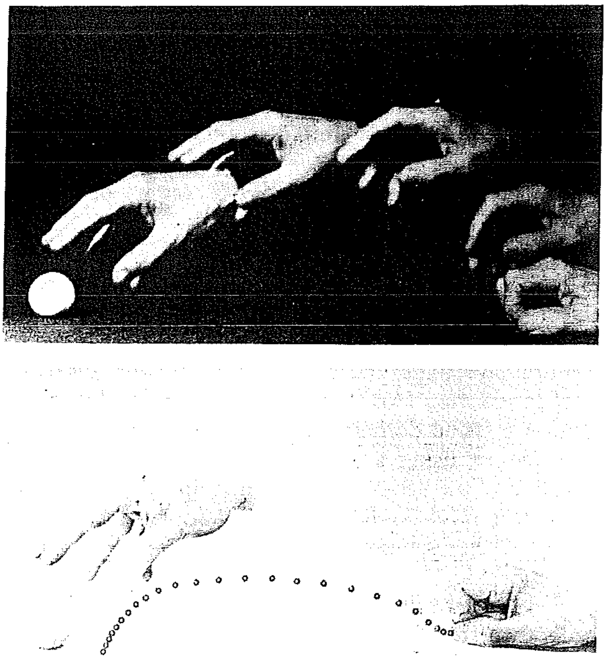

Figure 3. (Top) Superimposed view of hand shape and position in reaching for a ball. (Bottom) Circles indicate position of thumb-tip in successive movie frames and the separation into a slow phase followed by a fast phase (Jeannerod 1981).

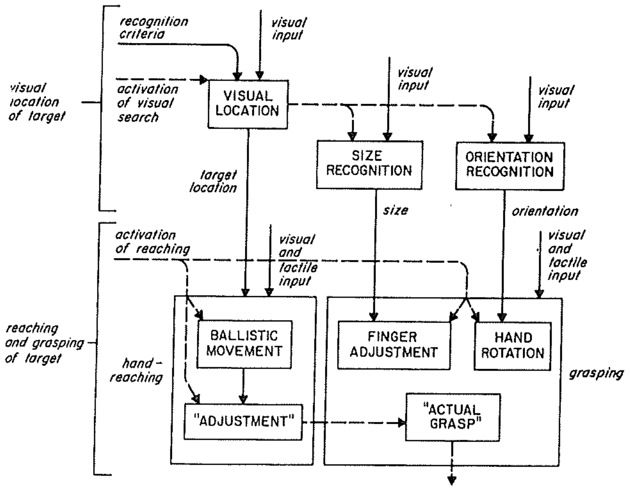

Figure 4. A coordinated control program for reaching toward and grasping an object (adapted from Arbib 1981).

Figure 4 shows a tentative program for this behavior. Solid lines indicate the transfer of data from one schema to another and dashed lines indicate the transfer of activation. (Appendix B describes a formalism for such coordinated control programs – there called "task units" – developed in the study of distributed control of dextrous robot hands.) The top half of the figure shows three perceptual schemas: Successfully locating the object activates schemas for recognizing its size and orientation. The outputs of these perceptual schemas are available for the control of the hand movement, which in turn involves the concurrent activation of two motor schemas. One of these moves the arm to transport the hand toward the object, the other concurrently preshapes the hand, with finger separation and orientation guided by the output of the appropriate perceptual schemas. We note that it is only the completion of the fast phase of hand transfer that triggers the slow phase of hand transfer and "wakes up" the final stage of the grasping schema to shape the fingers using tactile feedback. This simple example shows that a coordinated control program involves both concurrency and seriality of schema activation and it exemplifies how vision and touch interact. In some cases, a perceptual schema is an explicit representation of the environment on which planning can draw; in other cases, it is implicit, embedded in a motor schema and providing the tuning of parameters, or the setting of bounds for those parameters, directly. Finally, we have distinguished two types of motor control – a fast ballistic kind, and a slow passage to equilibrium. (For further work motivated by Figure 4, see Arbib et al. [1985] and Iberall et al. [1986], which introduce the notions of virtual finger and opposition space, respectively.)

We return now to our general consideration of animal behavior. As action continues, the current plan may continue to be executed simply with tuning or updating of parameters; or, because of some unexpected occurrence or completion of the current plan, some form of replanning or new planning may be required. From this viewpoint, both the schema assemblage and the current plan are temporally dynamic processes. They evolve with time, and so does the interaction between the two. Motor schemas change the environment and vice versa: Changes are remembered as well as anticipated. This provides the basis for learning. Here it seems useful to distinguish *episodic learning* (Squire's [1986] "declarative" memory), in which certain features of the schema assemblage and plan at a particular time are remembered together, from *skill learning* (Squire's "procedural" memory), in which a plan and its parameters are updated over time to better tune them to environmental circumstances, but without conscious recall of the learning experience (see Arbib [1985] and Arbib and Hesse [1986] for a first sketch of a schema-theoretic account of consciousness). The memory of particular episodes may be indexed with respect to schemas for general skills, and so in this way generalization may override specific memories.

Novel inputs (e.g., coming upon an unexpected obstacle) can alter the elaboration of high-level structures into lower-level tests and actions that in turn call upon the interaction of motor and sensory systems. We study programs that are part of the internal state of the system and can flexibly guide ongoing action in terms of internal goals or drives and external circumstances. Our thesis, then, is that the perception of an object (activating appropriate perceptual schemas) involves gaining access to routines for interaction with the object (motor schemas) but not necessarily the execution of even one of these routines. Although an animal may perceive many aspects of its environment, only a few of these can at any time become primary loci of interaction. Perception therefore "activates" (i.e., defines a search space; draws a map), and planning "concentrates" (lays out the route to be followed).

Our framework for analyzing visually guided behavior of a complex organism is thus based on these general premises.

1. As the organism moves – making, executing, and updating plans – it must maintain an up-to-date representation of its relationship with its environment.

2. The "model of the environment" is an active, information-seeking process composed of an *assemblage* of *perceptual schemas*, each instantiation of which represents a distinct domain of interaction with relevant properties, such as size and motion, represented by the current values of parameters of the schema.

3. The activation of perceptual schemas provides access to related *motor schemas* but does not necessarily entail their execution. *Planning* is required to determine the actual course of action. The plan is updated as action affords perceptual updating of the internal model.

4. The plan of action is to be thought of as a *coordinated control program* composed of motor schemas. Each motor schema is viewed as an adaptive controller that can update its representation of the object being controlled. Thus the adaptation procedure can be viewed as a perceptual schema embedded in a motor schema.

Behavior is as real as anatomy. However, when we spell out a network of interacting schemas that subserves it, we are dealing with theoretical constructs. Are schemas "real"? My provisional answer is that they are approximations to reality, just as many physicists would argue that concepts of the electron at the turn of the century were approximations to a reality that was revealed only with the development of quantum mechanics (a reality that is now itself seen to be but an approximation, as measured by the pragmatic criterion, as we probe the world of quarks and strings). In any particular diagram, such as Figure 4, the schemas may be too neatly separated. The schemas become "more real" as their functional analysis is refined into assemblages/programs of subschemas that allow either a more subtle analysis of behavior or an improved mapping of function to neural structure.

The quest is for a general theory of interacting perceptual and motor schemas adequate for top-down analysis of action-oriented perception or perceptually guided action, or for language (Arbib et al. 1987) and other cognitive processes. The general framework must be complemented by concepts specific to a particular domain. In the next three sections, we accordingly turn to the study of visuomotor cordination in the frog and toad to exemplify how a schema-theoretic analysis may be integrated with the search for the specific neural codes used by the various schemas and brain regions.

3. Schema models of *Rana computatrix*

Lettvin et al. (1959) initiated the behaviorally oriented study of the frog visual system with their classification of retinal ganglion cells into four classes, each projecting to a retinotopic map at a different depth in the optic tectum, the four maps in register. We view the analysis of such interactions between layers of neurons as a major approach to modeling "the style of the brain." In this section we present models of visuomotor coordination in the frog and toad at the level of schemas; the next two sections illustrate the general view of cooperative computation between neurons within a layer, and between layers in specific models exhibiting cooperative computation. We will make the following points:

1. A given schema, defined functionally, may be distributed across more than one brain region; conversely, a given brain region may be involved in many schemas. A top-down analysis may advance specific hypotheses about the localization of (sub)schemas in the brain, and these can be tested by lesion experiments.

2. In some cases we will then proceed to model each schema by interacting layers of neuronlike elements, or by nets of "intermediate-level" units. Even if the nets are little constrained by anatomy or physiology, such studies can be valuable in extending our understanding of "parallel distributed processing/connectionist" approaches to cognition and of the properties of neural networks, better preparing us to handle new data as they become available.

3. However, once a schema-theoretic model of some animal behavior has been refined to yield hypotheses about the localization of schemas, we can then model a brain region by seeing whether its known neural circuitry can indeed be shown to implement the posited schema. In some cases the model will suggest properties of the

circuitry that have not yet been tested, thus laying the ground for new experiments.

4. Once a number of models have been established, further modeling should be incremental, in that new models should refine, modify, and build upon prior models, rather than being constructed *ab initio*. Incremental modeling will not always work, however, and new data may lead to extensive reformulation. In a sense, the less often this happens, the better the modeling strategy.

3.1. Schemas for pattern recognition. Lettvin et al. (1959) found that group 2 retinal cells responded best to the movement of a small object within the receptive field and group 4 cells responded best to the passage of a large object across the receptive field. It became common to speak of R2 cells as "bug detectors" (Barlow [1953] referred to R3 ganglion cells as "fly detectors") and of R4 cells as "enemy detectors," although subsequent studies make it clear that a given frog or toad behavior will depend on far more than the activity of a single class of retinal ganglion cells (Ewert 1976).

Given the mapping of retinal "feature detectors" to the tectum and the fact that tectal stimulation could elicit a snapping response, it became commonplace to view one task of the tectum to be directing the snapping of the animal at small moving objects – it being known that the frog would jump away from large moving objects and would not respond when there were only stationary objects. This might suggest that the animal is controlled by, *inter alia*, two schemas, one for prey catching, which is triggered by the recognition of small moving objects, and one for predator avoidance, which is triggered by large moving objects. However, Ewert (1976) has observed (Figure 5) that animals with lesions of the pretectum will snap at large moving objects that a normal toad will avoid. This suggests a new analysis in terms of a prey-selection schema that can be activated by moving objects of any size, and a predator-recognition schema that serves not only to activate avoidance behavior but also to inhibit prey acquisition. Thus, even gross lesion studies can distinguish between alternative top-down analyses of a given behavior.

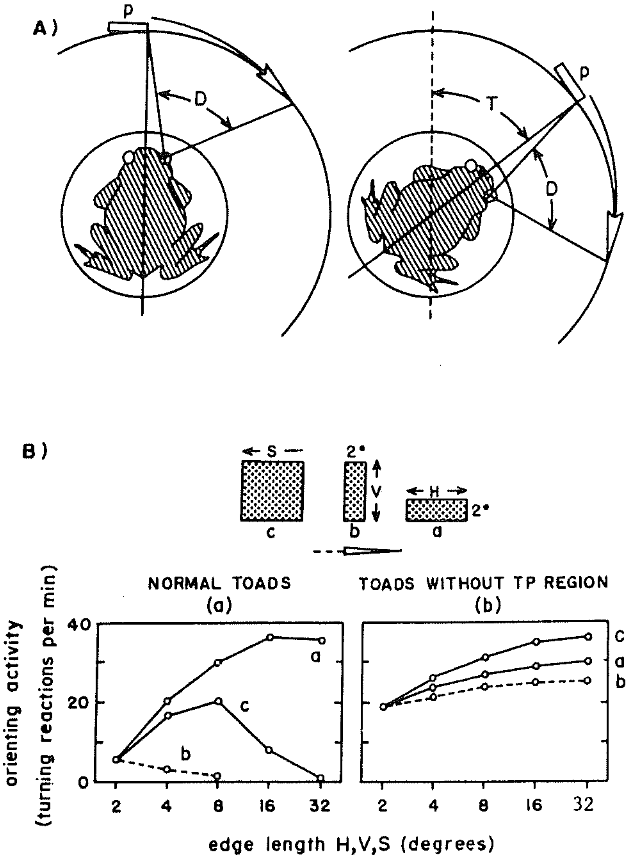

Figure 6. Prey-catching orienting behavior to different configurations of the stimulus. (A) Turning reaction to the stimulus presentation. Note that D, the effective angular displacement of the stimulus, differs from T, the angle of turning movements. (B) Orienting activity to three stimulus configurations, parallel ("worm": type a) and perpendicular ("antiworm": type b) rectangles, and squares (type c). (B.a) Normal animal's response becomes more frequent when we increase the dimension (H) of a stimulus of type a, whereas response frequency rapidly drops to zero when we increase the dimension (V) of a type b stimulus, and a sort of summation of these two responses is obtained when we increase both dimensions of stimulus type c. (B.b) This discrimination is lost in toads with pretectal lesions (from Ewert 1976).

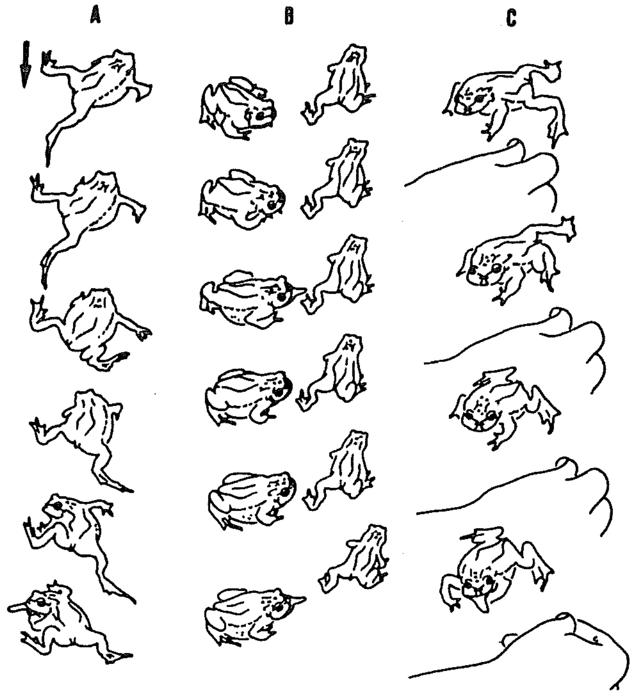

Figure 5. Three sequences showing a toad with a pretectal lesion snapping at moving objects at which the normal toad would not snap: (a) its own foot; (b) the foot of another toad; (c) experimenter's hand (Ewert 1976).

Of course, such an analysis can be refined by more detailed behavioral studies that let us determine what features of a moving object serve to elicit one form of behavior or another. For example, Ewert (Figure 6) placed a toad in a perspex cylinder from which it could see a stimulus object being rotated around it. He then observed how often the animal would respond with an orienting movement (this frequency being his measure of how "preylike" the object was) for different stimulus objects. As we can see from Figure 6B, the wormlike stimulus (rectangle moved in the direction of its long axis) proved increasingly effective with increasing length, whereas for 8° or more extension on its long axis, the antiworm stimulus (rectangle moved in the direction orthogonal to its long axis) proved ineffective in releasing orienting behavior. The square showed an intermediate behavior; the response it elicits rises to a maximum at 8°, but is extinguished by 32°. (See Ewert's accompanying target article, this issue.)

With such quantitative data at hand, Ewert and von Seelen (1974) produced the top-down model shown in Figure 7. As seen in Figure 7A, they postulated that retinal output was passed in parallel to a tectal "worm filter" and a thalamic "antiworm filter," with the output of the latter serving to inhibit tectal (type II) activity excited by the former. As we can see in Figure 7B, a worm stimulus would then tend to yield strong excitation of the worm filter, which would be little inhibited by the thalamic antiworm response, thus yielding a vigorous output; the antiworm (Figure 7D), on the other hand, would yield weak tectal type I activity, strong thalamic activity, and resultant weak tectal output. The square would yield intermediate behavior. Ewert and von Seelen were able to adjust the parameters in this model to fit the data over a linear subrange of the results. In suggesting how we might build on this, however, our main point is not that the model is restricted to linearity, but that it is "lumped" in both space and time in that it omits fine details in both these dimensions. That is, although the average rate of response of the output correlates well with the average turning rate of the toad, the model can neither explain the spatial locus at which the toad snaps nor the time at which it snaps. We shall thus be interested (Section 4) in the development of unlumped models, which can indeed explain the spatiotemporal distribution of the animal's behavior.

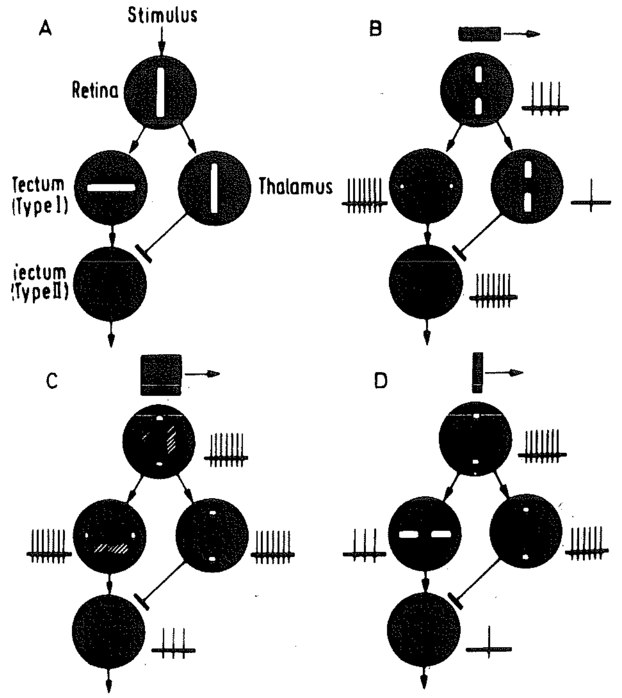

Figure 7. A lumped model of prey–predator discrimination: A worm filter provides excitatory input to an output cell, which also receives inhibitory input from an antiworm cell. The result is that (B) a worm provides a potent response; (C) a square provides an immediate response; and (D) an antiworm yields little or no response (after Ewert & Von Seelen 1974).

3.2. A model of prey selection. In much visually guided behavior the animal does not simply respond to a single stimulus but rather to some property of the overall configuration. Consider, for example, the snapping behavior of frogs confronted with one or more flylike stimulus. Ingle (1968) found that it is only in a restricted region around the head of a frog that the presence of a flylike stimulus elicits a snap; that is, the frog turns so that its midline is pointed at the stimulus and captures it with its tongue. There is a larger zone in which the frog only orients, and beyond that zone the stimulus elicits no response at all. When confronted with two "flies" within the snapping zone, either of which is vigorous enough that alone it could elicit a snapping response, the frog exhibits one of three reactions: It snaps at one of the flies, it does not snap at all, or it snaps in between at the "average fly." Didday (1976) offered the simple model of this choice behavior as shown in Figure 8a. It is presented here not to represent the state of the art, but rather as a

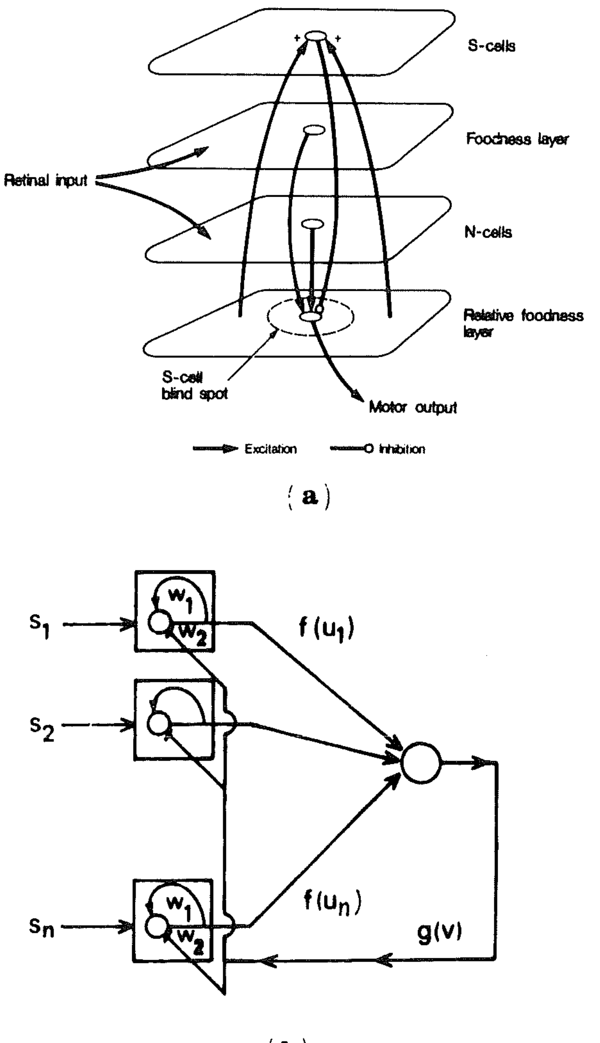

(a)

(b)

Figure 8. (a) Schematic view of Didday's model of interacting layers of neurons subserving prey selection. (b) Primitive cooperation model in which the layer of S-cells of (a) is replaced by a single inhibitory neuron (from Amari & Arbib 1977).

clear example of the distributed processing of structured stimuli. Didday used the term "foodness" to refer to the parameter representing the extent to which a stimulus could, when presented alone, elicit a snapping response. The task was to design a network that could take a position-tagged "foodness array" and ensure that usually only one region of activity would influence the motor control system. The model maintains the spatial distribution of information; new circuitry is introduced allowing different regions of the tectum to compete so that normally only the most active region provides an above-threshold input to the motor circuitry. To achieve this effect he first introduces a new layer of cells that is in retinotopic correspondence to the "foodness layer" and that yields the input to the motor circuitry. In some sense, then, it is to be "relative foodness" rather than foodness that describes the receptive field activity appropriate to a cell of this layer.

Didday's transformation scheme from foodness to relative-foodness uses a population of "S-cells" that are in topographic correspondence with the other layers. Each S-cell inhibits the activity that cells in its region of the relative-foodness layer receive from the corresponding cells in the foodness layer by an amount that increases with greater activity *outside* its particular region. This ensures that high activity in a region of the foodness layer penetrates only if the surrounding areas do not contain sufficiently high activity to block it. (Amari and Arbib [1977] present a functionally equivalent but more realistic model in which S-cells have no blind spot, but do receive recurrent local excitation.) When we examine the behavior of such a network, we find that plausible interconnection schemes yield the following properties:

1. If the activity in one region far exceeds the activity in any other region, then this region eventually overwhelms all other regions, and the animal snaps at the corresponding space.
2. If two regions have sufficiently similar activity levels then (a) they may both (providing they are very active) overwhelm the other regions and simultaneously take command, with the result that the frog snaps between the regions; or (b) the two active regions may simply turn down each other's activity, as well as activity in other regions, to the point that neither is sufficient to take command. In this case the frog remains immobile, ignoring the two "flies."

One trouble with the circuitry as so far described is that the buildup of inhibition on the S-cells precludes the system's quick response to new stimuli. If in case 2b above, for example, one of those two very active regions were suddenly to become more active, then the deadlock ought to be broken quickly. In the network so far described, however, the new activity cannot easily break through the inhibition built up on the S-cell in its region. In other words, there is hysteresis. Didday thus introduced an "N-cell" for each S-cell. The job of an N-cell is to monitor temporal changes in the activity of its region. Should it detect a sufficiently dramatic increase in the region's activity, it then overrides the inhibition on the S-cell and permits this new level of activity to enter the relative foodness layer. With this scheme the inertia of the old model is overcome, and the system can respond rapidly to significant new stimuli (see Amari and Arbib 1977 for a mathematical analysis). Didday hypothesized that the S-cells and N-cells modeled the "sameness" and "newness" cells, respectively, that had been observed in the frog tectum. Regrettably, no experiments have been done to test this hypothesis.

3.3. Schemas for depth and detours. We next describe a high-level model that addresses data on depth and detour behavior in toads in terms of interacting schemas. Collett (1982) has shown that a toad, confronted with a barrier beyond which a worm can be seen, may proceed directly toward the prey or it may sidestep the barrier and then approach the prey. In Figure 9, each dot–dash line pair indicates a position of the toad, with the dot indicating the position of its head and the line indicating its orientation, as seen from above. The whole sequence on the right-hand side of the figure represents one behavior of the animal in response to the stimulus shown below it. The row of dots represents a paling fence; the T, an opaque barrier; the circle with squiggles in it, a dish of mealworms. The mealworms are initially visible to the toad and lead, in combination with the fence, to the whole sequence of behavior in which the animal sidesteps around the fence, pauses, and then continues to the position marked "stop." What is worth noting is that soon after its initial movement the toad can no longer see the worms, and yet it proceeds along a trajectory whose final stage clearly indicates that the animal has retained an accurate representation of their position. However, the final approach is aborted by the lack of adequate stimuli.

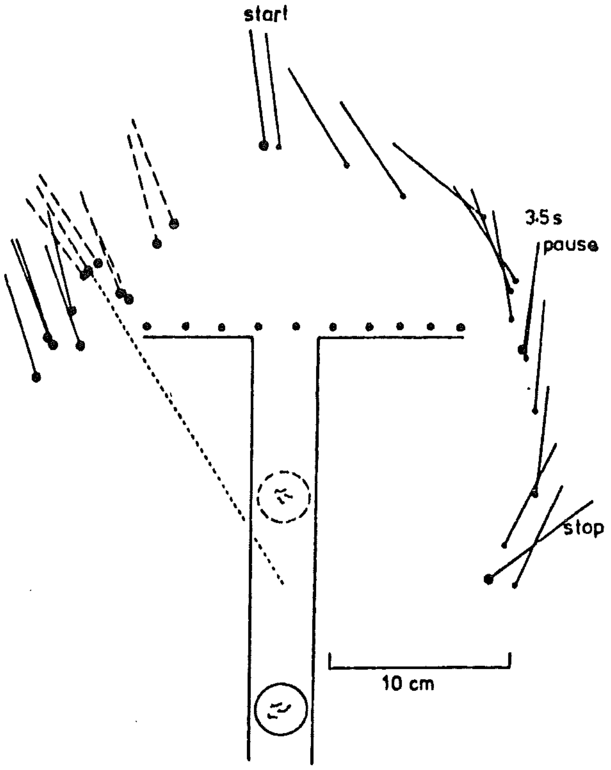

Figure 9. A trajectory of a toad that has sighted a number of mealworms behind a paling fence (the row of dots) and then detours around the fence to approach the prey. Note that when it stops (the prey no longer being in view) its position shows that it has retained a representation of prey position, which is relatively accurate despite the intervening movement of the toad (Collett 1982).

Epstein (1979) adapted Didday's simple model of the tectum as a row of neurons selecting its maximal stimulus by positing that each visible preylike stimulus provides a tectal input with a sharp peak at the tectal location corresponding retinotopically to the position of the stimulus in the visual field, with an exponential decay away from the peak. A barrier, on the other hand, provides a trough of inhibition whose tectal extent is slightly greater, retinotopically, than the extent of the barrier in the visual field. Epstein's model can exhibit choice of a target in the direction of the prey or the barrier edge, but not the spatial structure of the behavior.

Given that the behavior of the toad – whether approaching the prey directly or detouring around the barrier – depends on how far behind the barrier the worms are, a full model of this behavior must incorporate an analysis of the animal's perception of depth. To address this, Arbib and House (1987) gave two models for detour behavior that make use of separate depth maps for prey and barriers. In the first, the Orientation Model (Figure 10) the retinal output of both eyes is processed for "barrier" and "worm" recognition to provide separate depth mappings for barrier and worm. We suggest that the animal's behavior reflects the combined effects of prey "attraction" and barrier "repulsion." Formally generalizing Epstein's model, the barrier map B is convolved with a mask I which provides a (position-dependent) inhibitory effect for each fencepost, whereas the worm depth map W is convolved with a mask E that provides an excitatory effect for each worm. The resultant map

$$T = B{*}I + W{*}E$$

is then subject to further processing, which will determine the chosen target. E is an excitatory mask that projects broadly laterally and somewhat less broadly toward the animal; I is an inhibitory mask with a short distance behind the edge where there is little inhibition (beyond this, inhibition is equally strong at all distances). The total excitation, T, is summed in each direction, and then a maximum selector network chooses the direction with maximal activity. If this corresponds to the prey, the animal will approach and snap; otherwise, further processing is required. We postulate that each component of the detour behavior (sidestep, orient, snap, etc.) is governed by a specific motor schema. Ingle (1982) offers some clues as to their localization: He finds that a lesion of the crossed tectofugal pathway will remove orienting; lesioning the crossed pretectofugal pathway will block sidestepping; and lesions of the uncrossed tectofugal pathway will block snapping.

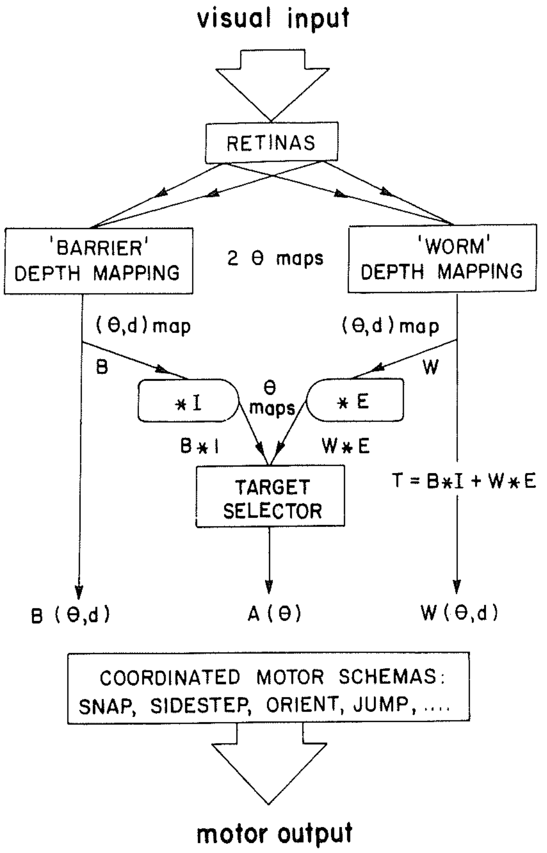

Figure 10. Conceptual schematic of visual/motor pathway: Assumptions made in this diagram are (1) that separate depth maps are maintained for prey and barrier stimuli, (2) direction for an orientation turn is obtained by combining information from these two depth maps, and (3) information on preferred orientation and depth of prey and barriers is available simultaneously to motor schemas. These schemas are capable of integrating this information to produce a coordinated motor output (Arbib & House 1987).

In their second model, the Path-Planning Model, Arbib and House (1984) associate with each point of the depth map a two-dimensional vector. In place of a single scalar indicating a measure of confidence that there is a target for the first move at the corresponding position in the visual field, the vector is to indicate the preferred direction in which the animal should move were it to find itself at the corresponding position. The model specifies how this vector field is generated and begins to specify how the vector field is processed to determine the appropriate parameters for the coordinated activation of motor schemas. Each prey object sets up an attractant field, while each fencepost sets up a field for a predominantly lateral movement relative to the position of the post from the viewpoint of the animal. Arbib and House suggest that in the case of a "tracking creature" like the gerbil, the vector field is integrated to yield a variety of trajectories, with a weight factor for each trajectory, whereas, in a "ballistic creature" like the frog or toad, processing yields a map of motor targets, appropriately labeled as to type. The current model uses vectors encoding components of forward and lateral motion; future work will explore the hypothesis that a particular vector would have components governing sidestepping, turning, and snapping. It is an open question whether the components of the vector would be expressed in adjacent nerve cells or distributed across different regions of the brain.

3.4. Schemas for prey acquisition. Lara et al. (1984) offer an alternative model of detour behavior in the presence of barriers with gaps in which the recognition of gaps is an explicit step in detour computation. The same paper also offers models – at the level of interacting schemas rather than layers of neuronlike elements – for prey acquisition in environments containing chasms as well as barriers, and for predator avoidance. We now turn to a presentation of the schemas for prey acquisition by toads posited in this approach. In what follows, the figures use the

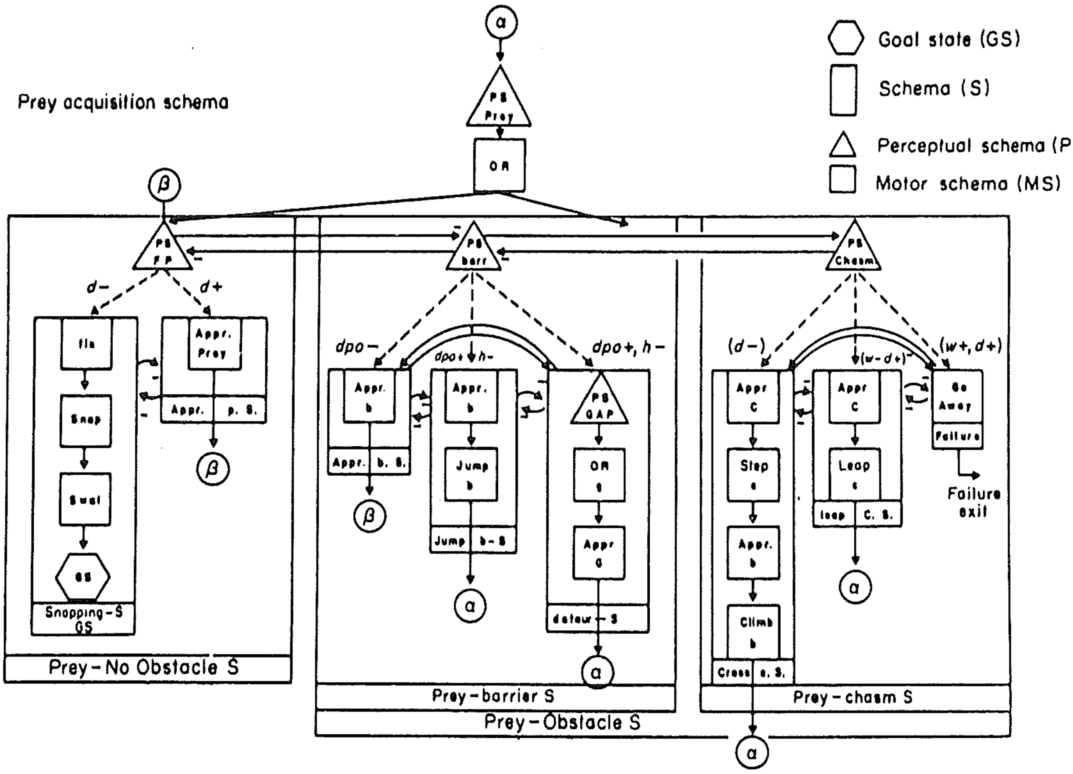

Figure 11. If prey is in the visual field, the orientational motor schema *OR* is activated. The consequent unfolding of action depends on whether a barrier (activating perceptual schema *barr*) or a chasm (activating perceptual schema *chasm*) or "free space" (activating perceptual schema *FP*) is perceived. See text for further details (from Lara et al. 1984).

abbreviations PS, MS, S, and GS for perceptual schema, motor schema, schema, and goal state, respectively.

The prey-acquisition schema (Figure 11) is activated by the instantiation of the perceptual schema for *prey,* but the actual motor schemas executed will depend on whether further perceptual processing activates the "free prey" perceptual schema *FP* representing a situation in which no obstacles intervene between the animal and its prey; perceptual schema *barr,* which represents a situation in which a barrier intervenes; and perceptual schema *chasm,* which represents an intervening barrier. The activation of the corresponding schema – *prey-no obstacle, prey-barrier,* or *prey-chasm,* respectively – is represented in Figure 11 as the outcome of competitive interaction between the three perceptual schemas.

If perceptual schema *FP* is instantiated, it generates a parameter *d* for the distance of the prey. If *d* is small (d−), then the animal fixates, snaps, and swallows the prey to achieve the goal state. If *d* is larger (d+), then the animal approaches the prey. The arrow to β indicates that the animal will return to the circle marked β above – so long as the prey remains visible, the animal will continue to approach it until *d* is small enough for it to snap and swallow the prey. It might be argued that both of the β returns in the figure should be replaced by returns to α – since a barrier might be interposed or become visible after the animal has executed an approach motor schema. In any case, such returns (to be discussed further below) represent the behavioral fact that the toad seems to execute a single behavior (which may itself be composite, as in fixate→snap→swallow) to completion, with the passage to the next requiring a fresh perceptual trigger, rather than formulating some overall plan to be pursued and modulated in the light of changing circumstances (in the way that a human, going to get some object, would initiate a search strategy should the object not be in its expected place).

If perceptual schema *barr* is instantiated, it makes available two parameters: *dpo,* the distance of the prey behind the barrier, and *h,* the height of the barrier. If *dpo* is small or negative (dpo−), the toad approaches the barrier and proceeds as if no barrier were present (transfer to β). If the prey is further behind the barrier and the barrier is low (dpo+, h−), the toad will approach and jump over the barrier, whereas if the barrier is high (dpo+, h+), the animal will make a detour. More specifically, instantiation of the perceptual schema for a gap will elicit the motor schema for orienting to (*OR g*) and approaching (*Appr g*) that gap. It seems necessary to include the "gap" at the end of a barrier as also activating the gap perceptual schema. In either dpo+ case, control is returned to α after execution of the specified motor schema.

Finally, if perceptual schema *chasm* is activated, it makes available parameters representing the depth *d* and width *w* of the chasm. If the depth is small (d−) the toad will walk across the floor of the chasm; if the chasm is deep but not wide (w−, d+), the toad will leap the chasm; but if the chasm is wide and deep (w+, d+), the animal will not approach the chasm.

The general considerations of Section 2, which influenced the Arbib-House model of detour behavior, suggest that the animal's behavior should depend on the representation of the environment by an assemblage of schemas embodying the spatial relations between multiple objects in the environment, rather than sequential activation of a perceptual schema for one prey object and then at most one obstacle. In fact, in the description Lara et al. (1984) give of the actual implementation of approaching gaps, they do use a mode of competitive interaction, which is only hinted at in Figure 11, by the various arrows labeled "-" to express competitive interactions. More generally, the idea seems to be that the gap schema is instantiated for each gap in the environment; each schema is given an activity level based on its position relative to toad, prey, and other gaps; and it is the most active gap schema that provides the parameters for the activation of the next motor schema. Once this motor schema has been executed, the animal's behavior is determined *ab initio* by the activation of its perceptual schemas in its new situation.

Note that this explicit view of motor schema activation by the "winner" (cf. the Didday model of prey selection above) of perceptual schema competition obviates the α–β question raised above, since we simply postulate that, once triggered, a motor schema is executed to completion with the parameters supplied by the perceptual schemas, whereupon activation of perceptual schemas is reinitiated to trigger the next motor schema. As noted above, this seems an appropriate model for the frog and toad – the schema assemblage is completely "refreshed" after each integral action. This is in contrast to the action-perception cycle in human behavior, for example, in which activation of perceptual schemas serves to update an existent schema assemblage and coordinated control program, rather than to create new ones.

An example of two trajectories predicted by the model of Lara et al. is presented in Figure 12. In (a), the "toad" chooses the gap nearer the prey, but in (b), the presence of the wide gap in the further fence biases the animal's "choice" to favor the other gap in the nearer fence. Clearly, such predictions can be used to test the model, and data culled from observations of animal behavior can be used to refine the algorithmic specification of the constituent schemas, just as psychophysical data can be used to tune models of visual function in humans. Such specifications can serve as the endpoint of analysis for the cognitive scientist or ethologist who is not interested in neural mechanisms; but for the neuroscientist, they can provide the formal "specification" of the task whose implementation in neural circuitry is to be analyzed.

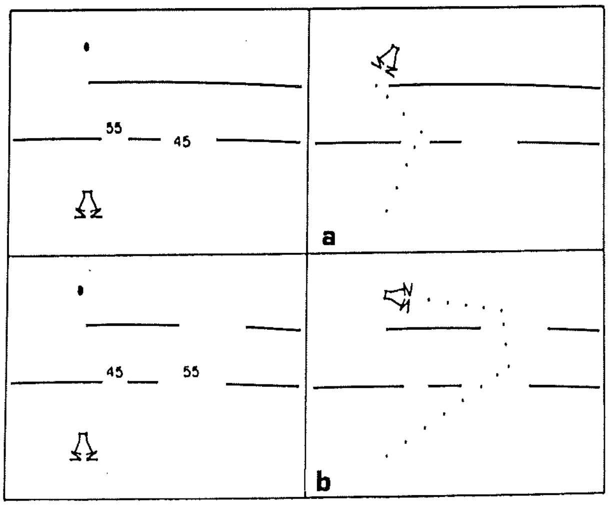

Figure 12. Computer simulation, using a program elaborating the schemas of Figure 11, of a toad's behavior in response to a worm in the presence of barriers. In (a) the toad always prefers the closer gap; in (b) the farther gap is wider and is the one chosen (from Lara et al. 1984).

4. Tectal column models of *Rana computatrix*

In this section and the next we proceed down a level from schemas, examining how certain of the schemas of the previous section can be captured in the cooperative computation of layers in anatomically defined regions of the anuran brain. In the present section, we will pursue the analysis yet another level downward, showing how the properties of certain layers can be implemented by models of neural networks constrained by neurophysiological data.

4.1. Facilitation of prey-catching behavior. Frogs and toads take a surprisingly long time to respond to a worm. Presenting a worm to a frog for 0.3 sec may yield no response, whereas orientation is highly likely to result from a 0.6 sec presentation. Ingle (1975) observed a facilitation effect: If a worm is presented initially for 0.3 sec, then removed, and then restored for only 0.3 sec, the second presentation suffices to elicit a response so long as the intervening delay is at most a few seconds. Ingle observed tectal cells whose time course of firing accorded well with this facilitation effect (Figure 13d). This leads us to a model (Lara et al. 1982) in which the "short-term memory" is encoded as reverberatory neural activity rather than as the short-term plastic changes in synaptic efficacy demonstrated, for example, by Kandel (e.g., 1978) in *Aplysia*. Our model is by no means the simplest model of facilitation – rather, it provides a reverberatory mechanism for facilitation consistent with Ingle's neurophysiology and the then known local neuroanatomy of the tectum.

The present model addressed facilitation at a single locus of the tectum. Interactions among a number of columns will be discussed below. The tectal columnar model is abstracted somewhat crudely from the anatomy of Székely and Lázár (1976) – and thus must be integrated with new and detailed anatomy and neurophysiology (e.g., Lázár et al. 1983). Each column (Figure 2d) comprises one pyramidal cell (PY) as sole output cell, one large pear-shaped cell (LP), one small pear-shaped cell (SP), and one stellate interneuron (SN). (The simulation results of Figures 13 and 14 were actually based on a larger column, with 1 PY, 3 LP, 2 SP, and 2 SN, but the results for the column of Figure 2d are essentially the same.) All cells are modeled as excitatory, save for the stellates. The retinal input to the model is a "foodness" measure, and activates the column through glomeruli with the dendrites of the LP cell. LP axons return to the glomerulus, providing a positive feedback loop. A branch of LP axons also goes to the SN cell. There is thus competition between "runaway positive feedback" and the stellate inhibition. (The equations used in the simulation can be found in Appendix 1 of Lara et al. 1982.)

To return to the tectal model: Glomerular activity also

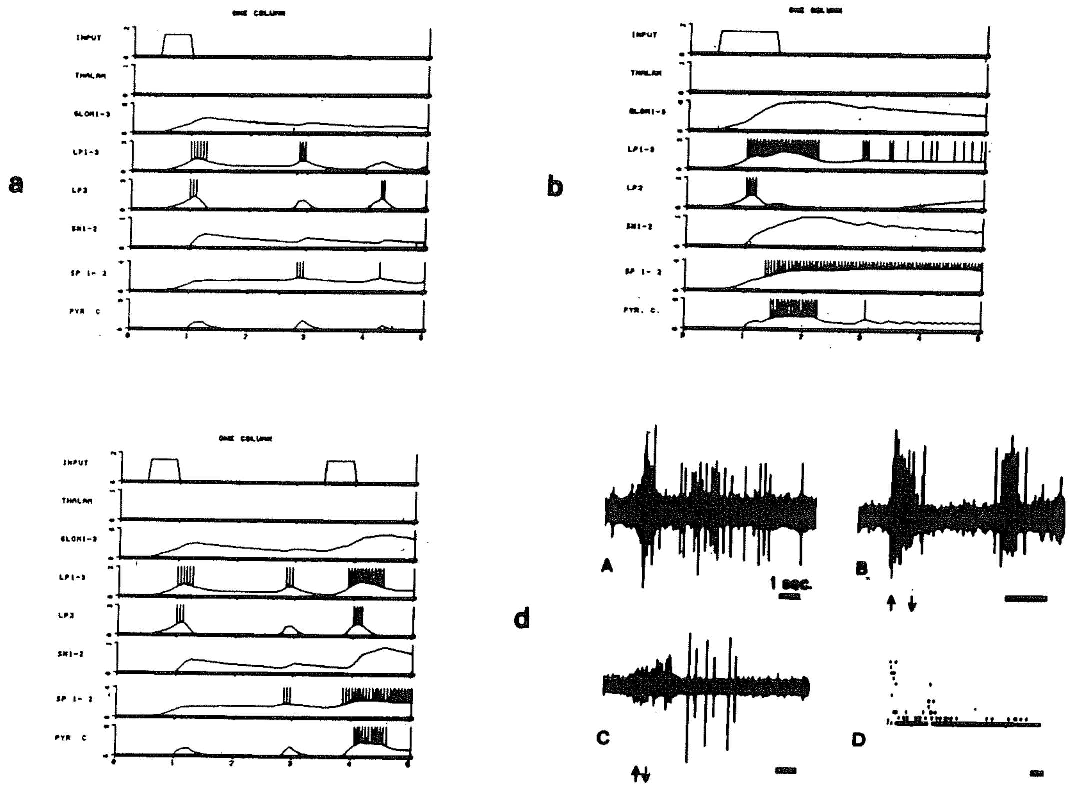

Figure 13. (a) Computer simulation of tectal cells' responses when a brief stimulus is presented. The onset of the stimulus produces a long-lasting depolarization in the glomerulus, which then fires the large pear-shaped cell (LP). This neuron in turn sends recurrent axons to the glomerulus and the stellate cell (SN), which acts as the inhibitory neuron in the column. When the inhibitory effect of SN releases the LP cell, a rebounding excitation occurs. The small pear-shaped cell is integrating the activity of GL, LP, and SN neurons to give a delayed short response. (b) If in the above situation we present a stimulus of longer duration, then we show that now the pyramidal neuron fires. In (c) we show that when a second stimulus of the "subthreshold duration" used in (a) is presented, the pyramidal cell (PY) responds. (The frequencies of the spikes are a graphical convention. The spikes are drawn simply to highlight when the membrane potential of a cell is above threshold.) (From Lara et al. 1982) (d) Physiological behavior of cells related to prey-catching facilitation: (A) A brief class 2 burst is followed by a delayed response of a tectal cell; (B) a tectal cell responds to the presentation of the stimulus and again with a delay; (C) a tectal neuron produces a delayed response to the presentation of the stimulus. (D) the poststimulus histogram of a tectal cell shows a delayed peak at 3 to 4 seconds (from Ingle 1975).

excites SP, which likewise sends its axon back to the glomerulus. SP also excites LP to recruit the activity of the column. PY is excited by both SP and LP. Clearly, the overall dynamics will depend upon the actual choice of excitatory and inhibitory weights and of membrane time constants to ensure that excitation of the input does not lead to runaway reverberation between the LP and its glomerulus, and that this activity is "chopped" by stellate inhibition to yield a period of alternating LP and SN activity. SP has a longer time constant and is recruited only if this alternating activity continues long enough. It required considerable computer experimentation to find weights that yield the neural patterns discussed below. More recently, Cervantes-Perez (1985) has given a mathematical analysis of how weighting patterns affect overall behavior. We hope our hypotheses concerning the ranges of the parameters involved in the model will stimulate more detailed anatomical and physiological studies of tectal activity.

In one simulation experiment, we graphed the activity of the pyramidal cell as a function of how long a single stimulus is applied (Figure 14a). There is, as in the experimental data, a critical presentation length below which there is no pyramidal response. Input activity activates LP, which reexcites the glomerulus but also excites the SN, which reduces LP activity. But if input continues, it builds on a larger base of glomerular activity, and so over time there is a buildup of LP-SN alternating firing. If the input is removed too soon, the reverberation will die out without activating SP enough for its activity to combine with LP activity and trigger the pyramidal output. If input is maintained long enough, the reverberation may continue, though not at a high enough level to trigger output. However, reintroducing input shortly after this "subthreshold" input can indeed "ride upon" the residual activity to build up to pyramidal output after a presentation time too short to yield output with an initial presentation.

4.2. Worm–antiworm discrimination. We will now try to explicate Ewert's behavioral data on prey–predator discrimination (Figure 6) in terms of neural networks at the level of detail of the tectal column. In doing this, we are helped by Ewert's recordings of the activity of different types of neurons in toads conducting the task of Figure 6. In Figure 15a, we see that TH3 neurons (in the thalamus/pretectum) have a response that is uniform for worms, increases with increasing length for antiworms, but is

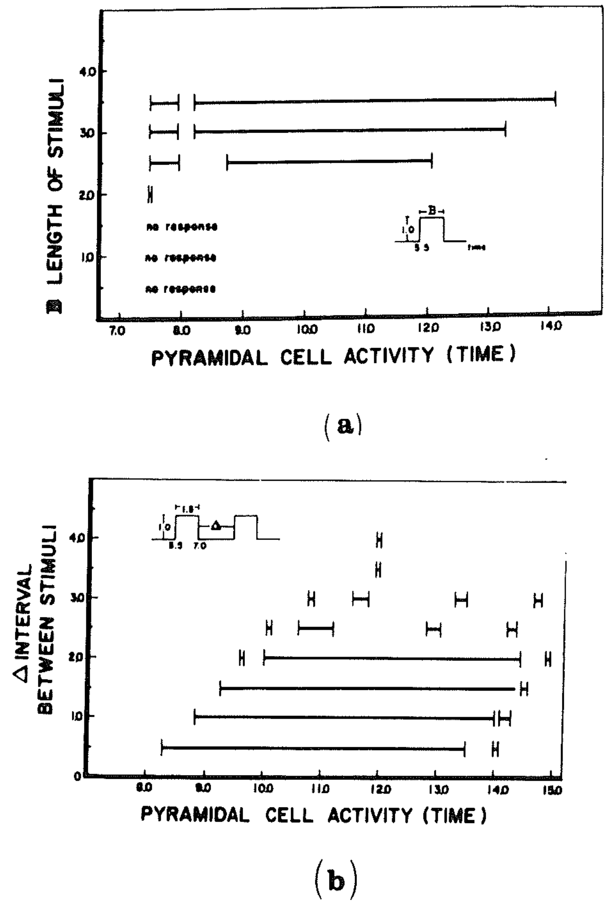

(a)

(b)

Figure 14. (a) Computer simulation of the PY behavior when stimuli are presented for different intervals. For each value B of the duration of the stimulus, we present a horizontal line showing the time during which the pyramidal cell is active. (b) Computer simulation of the temporal pattern of the facilitation process after the presentation of a brief stimulus. Here we show the period of pyramidal cell activity for various values Δ of the interval between a pair of stimuli each having duration B = 1.5 (Lara et al. 1982).

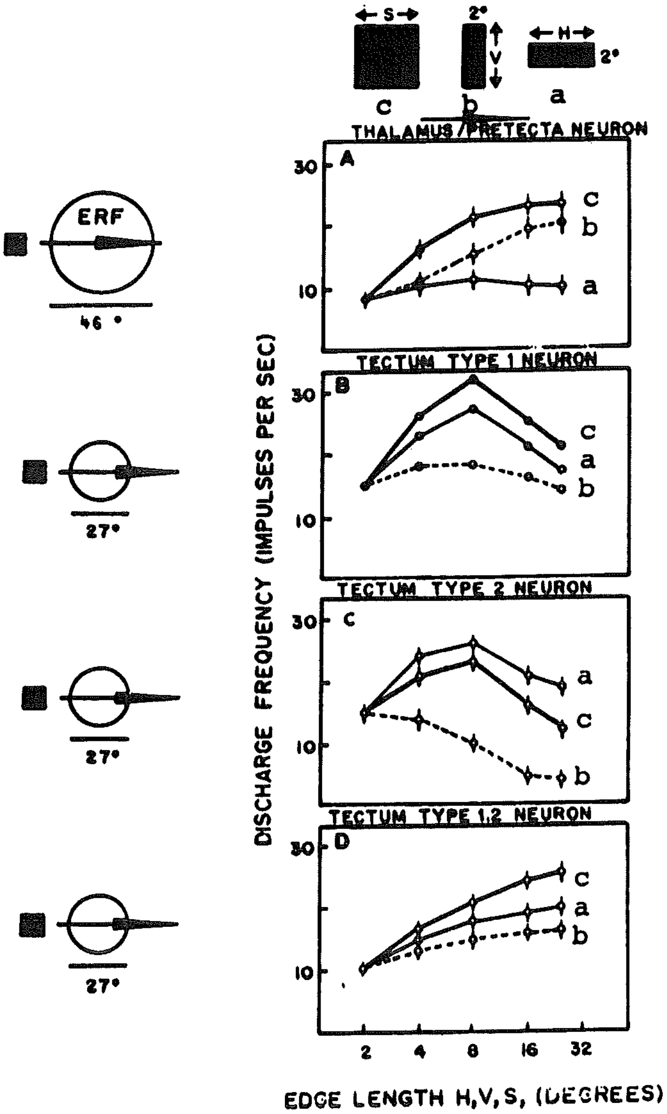

Figure 15. Tectal and pretectal cell response from common toads to different configurations of moving stimuli (see Figure 6 legend). (A) Response of a pretectal neuron TH3, which is mostly sensitive to large (type c) and perpendicular (type b) stimuli. (B) Response of a tectal cell T5(1), which is most sensitive to stimuli type c, then type a, and then type b. (C) Response of tectal neuron T5(2), which mostly prefers stimuli type a, then type c, and gives a very weak response to type b. This neuron's response resembles the animal's behavior. (D) Response of both tectal cells, T5(1) and T5(2), after thalamic pretectal lesions. It shows how the discriminative abilities of these cells are lost (from Ewert 1976).

greatest for squares, increasing with their size. The activity of tectum type I neurons (T5(1)) seems to correlate fairly well with increases in the length of the stimulus in the direction of motion, although it is somewhat higher for squares than for worms. It is the tectum type II neuron, T5(2), whose overall rate of *neural* response seems to best match the overall frequency of the *behavioral* response; and it is this averaged neural response that is explained by the Ewert–von Seelen model of Figure 7. We note, too, that Figure 6b shows that the discrimination among worm, antiworm, and square is essentially abolished by pretectal lesion. Our task now is to give a distributed model. As already discussed in Section 1, we want to do this in an incremental way – rather than seeking a minimal model of tectum to explain the specific pattern recognition task at hand, we try to model the tectum in terms of the basic columnar structure already presented.

The structure of our model is shown in Figure 16 (Cervantes-Perez 1985). Rather than using a single column to represent the tectum, we now use an 8×8 array of columns, and we provide the column not only with the R2 input of our initial study of facilitation, but also with R3 and R4 input. In addition, we represent the pretectum by an array of TH3 cells receiving R3 and R4 input. The retinal input is based on the ganglion cell response curves of Figure 17, in which only the average rate of firing of a cell is given for each stimulus, rather than the temporal pattern of that response. (Since the tectal response must depend on the spatiotemporal pattern of retinal ganglion cell firing, current work in my laboratory [Lee 1986] is aimed at more detailed modeling of the response of frog retina to varied activity.) The connections of R3 and R4 to TH3 are tuned to yield the responses shown in Figure 15A.

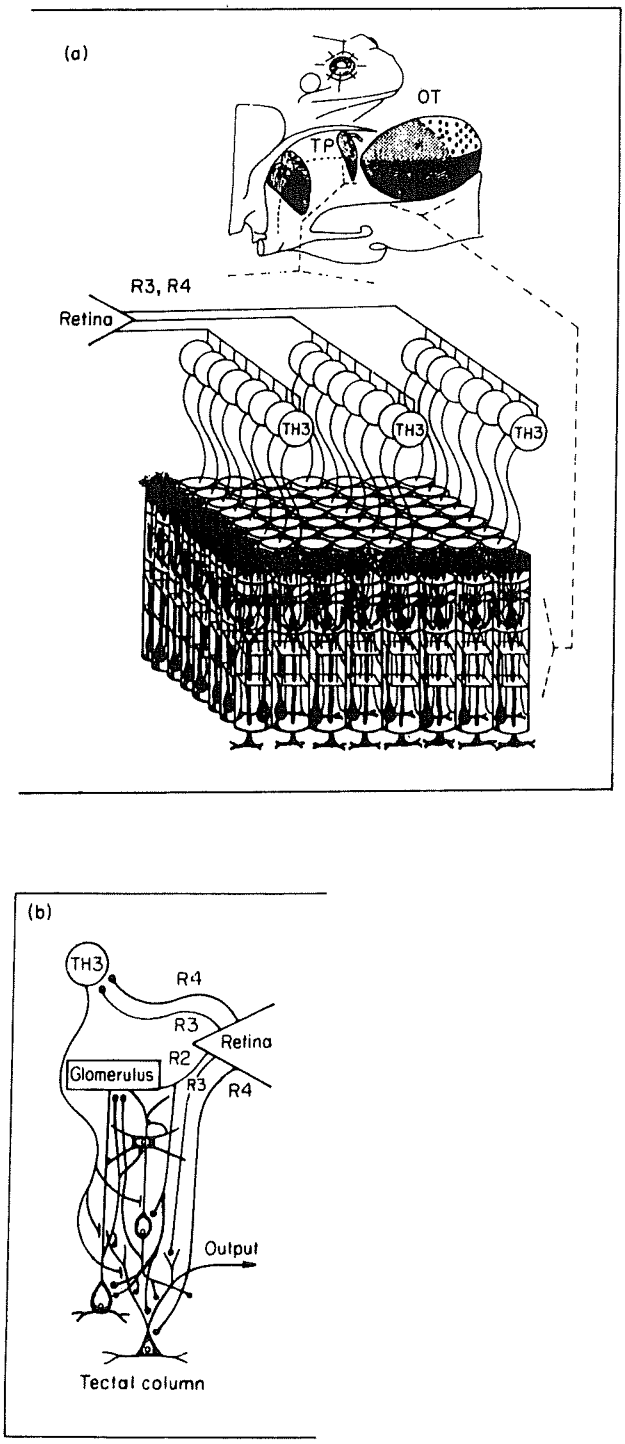

Figure 16. Interactions among retina, optic tectum, and pretectum. The retina sends fibers in a retinotopical fashion to both optic tectum (class R2, R3, and R4) and pretectum (class R3 and R4). (a) TH3 neurons also project retinotopically to the optic tectum. For simplicity we show only the projection of three rows of TH3 cells projecting upon the tectal columns. (b) A closer look at the interactions among retinal, tectal, and pretectal cells. The TH3 cell of the pretectal column inhibits LP, SP, and PY of the tectal column corresponding to its retinotopic projection (from Cervantes-Perez et al. 1985).

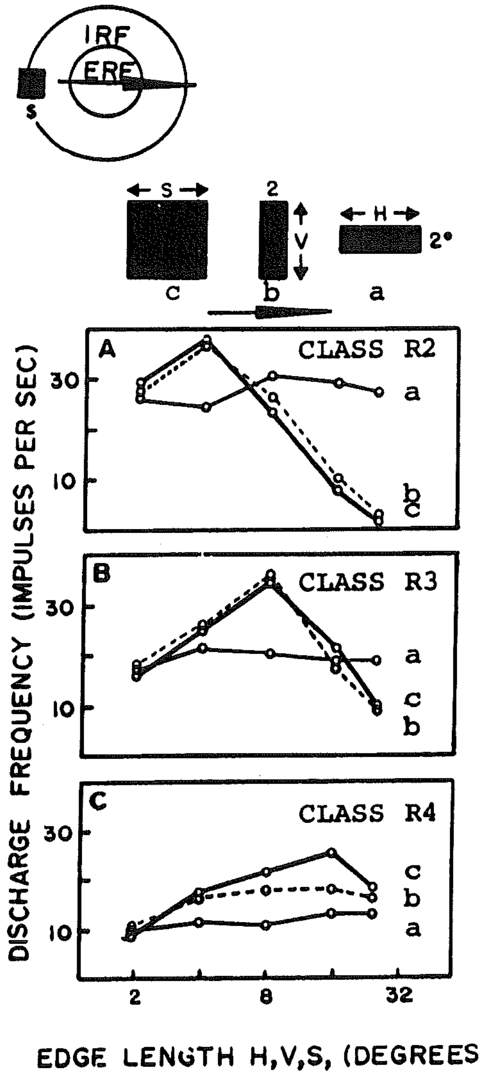

Figure 17. Retinal ganglion cells' response (classes R2, R3, and R4) to different configurations (type a, b, and c) of moving stimuli with a visual angular velocity of 7.6 degrees/sec. Left, they present a receptive field formed by a central excitatory (ERF) and a peripheral inhibitory (IRF) area. The three ganglion neurons respond almost with the same intensity to type a stimuli of different sizes. For stimuli of types b and c, ganglion cells R2 and R3 increase their rate of response up to their respective receptive field sizes, and then it drops down, whereas R4 increase their rate of response when the size of the stimulus increases, giving the strongest response to stimuli of type c (from Ewert 1976).

With appropriate setting of parameters, the model does indeed exhibit, in computer simulation, responses to moving stimuli of different types (Figure 18) that match well the neural data of Figure 15. However, it is clear that the model is only approximate at a quantitative level and that – if our goal is prediction of detailed neural firing rather than just a general understanding of pattern recognition networks – further work must be done on tuning the model parameters. Figures 19 and 20 give us a more detailed look at the output of the computer runs involved in generating the single points of the overall summation of the data in Figure 18. Figure 19 shows the response, for a brief period of time, of the 8×8 array of tectal columns to a moving worm (a), antiworm (b), and square (c) when *no*

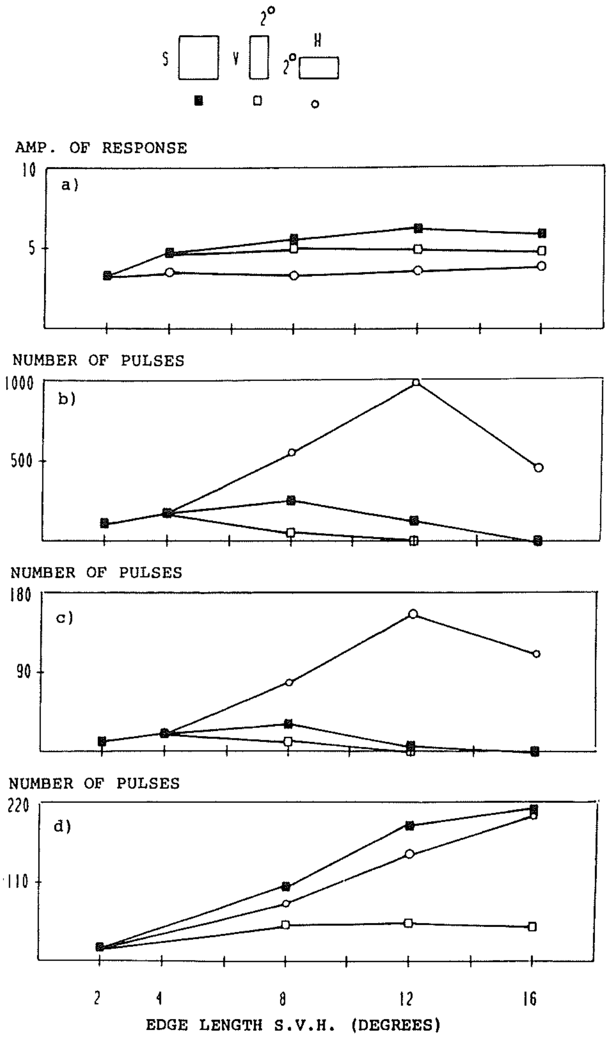

Figure 18. Computer simulation of the response of pretectal and tectal cells to different configurations (type a, b, and c) of moving stimuli (visual angular velocity of 8 degrees/sec). (a) Pretectal TH3 cell response: It is mostly sensitive to stimuli of type c and b. (b) Overall response of the tectum to the three types of stimulus (a, b, and c): Tectum response is mostly sensitive to stimulus type a and weaker to stimulus type b. (c) Response of a PY cell to the three different stimuli: It responds better to stimulus type a, then to type c, and it gives a very weak response to a stimulus of type b. This curve is equivalent to b. (d) PY response when pretectal ablation occurs: These cells are mostly sensitive to stimuli of type c or a, and less to those of type b (from Cervantes-Perez et al. 1985).

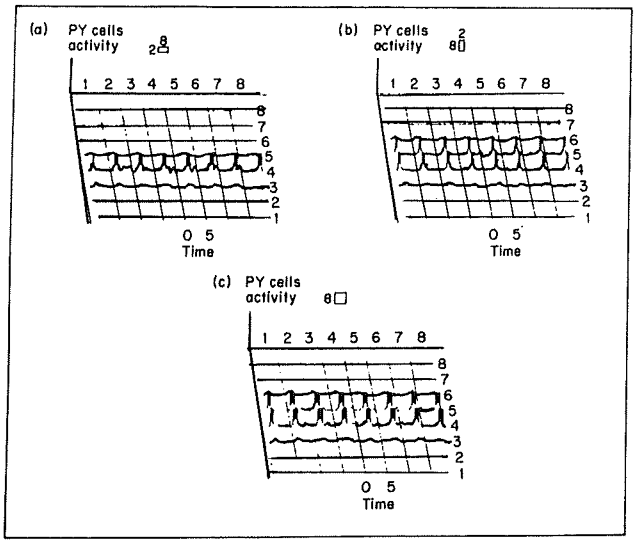

Figure 19. A computer simulation of the PY cell response of the 64 columns of the tectum to the different configurations of moving stimuli when pretectal ablation occurs: All stimuli are moved from left to right with a visual angular velocity of 8 degrees/sec. The figure shows four-dimensional graphs, where the x and y axes are used to represent the spatial localization of the (i, j) column. The y axis of this plane is also used to show the time scale for the response of every column's PY cell, and the vertical axis (z axis) represents its local membrane potential. When the PY local potential is above the threshold, this is indicated by spikes in the graphs. (a) Response to a "wormlike" stimulus (type a) of 8 × 2 degrees; (b) response to the same stimulus moved as "antiwormlike" (type b); (c) response to a square stimulus (type c) of 8 × 8 degrees. At the level of one PY cell, the tectal response is strongest to stimulus type c, then to type a, and finally to type b. The overall response of the tectum is also stronger to a type c stimulus. It is also more widespread, whereas the response to a "wormlike" stimulus, although likewise strong, is concentrated in a narrower area (from Cervantes-Perez et al. 1985).

pretectal influence is included in the model. Each figure represents an 8×8 array of graphs of neural activity corresponding to the pyramidal cell activity of the column, with time measured from left to right in each subgraph. We thus see that the response to a worm moving from left to right is delayed later and later as we look at tectal columns arrayed from left to right in the path corresponding retinotopically to the path of the worm. It is interesting to note that the response of individual cells to an antiworm is weaker than their response to a worm, but that, due to the lateral extent of the antiworm, more cells are excited in its passage – as is also the case for the square, which likewise yields more vigorous response on a cell-by-cell basis. Finally, Figure 20 shows the response of the full model in which the tectum is subject to pretectal inhibition. Here we see that there is a vigorous response to the worm, virtually no response to the antiworm, and only a weak response to the square. For further details of the model, the results of further simulation, a demonstration of directional invariance of response, and a discussion of motivation, the reader is referred to Cervantes-Perez et al. (1985).

4.3. Incremental modeling. We have demonstrated *incremental modeling* by showing how our model of a "tectal column" was introduced to explain certain facilitation effects in prey-catching behavior. In subsequent papers (not discussed here), a linear array of such columns has been used to model certain data on size dependence of prey-catching activity in toads; inhibition from pretectum to such an array was then introduced to model the behavior of an animal confronted with more than one prey stimulus. We have just discussed a fourth stage, modeling of the tectum as an 8×8 array of tectal columns in interaction with retina and pretectum. Lara and Arbib (1985) have further developed the model to address certain issues in stimulus-specific habituation. These models form five stages in an "evolutionary" sequence for *Rana computatrix*, and incorporate models of the retina,

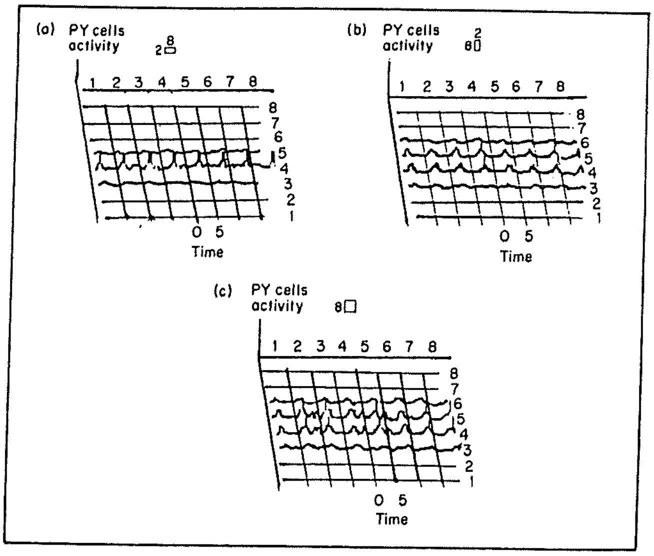

Figure 20. A computer simulation of the PY cell response of the 64 columns of the model of the interactions among retina, optic tectum, and pretectum. All stimuli are moved from left to right with a speed of 8 degrees/sec. (See Figure 19 legend for an explanation of the graph characteristics.) (a) Response to a "wormlike" stimulus of 8 × 2 degrees (type a); (b) response to an "antiwormlike" stimulus of 2 × 8 degrees (type b); and (c) response to a square stimulus of 8 × 8 degrees (type c). The response of the PY cells is strongest to stimulus type a, then to type c, and finally to type b. It can be seen that the response to a type c stimulus, although more widespread, is weaker than that to type a, whereas the weakest response is again to stimulus type b. PY cells are most sensitive to "wormlike" stimuli, rather than to squares or "antiwormlike" stimuli (from Cervantes-Perez et al. 1985).

tectum, and pretectum, as well as a forebrain region analogous to the hippocampus. Moreover, the models are designed to contribute to a "tool kit" that can help us study other schemas and neural systems (apart from the frog's and toad's) and that can encourage the comparative study of different systems. For example, the role of the stellate neuron in our tectum model is reminiscent of Purkinje inhibition of the positive feedback between cerebellar nuclei and reticular nuclei, a basic component of Boylls's (1975; 1976) model of cerebellar modulation of motor schemas, based on Tsukahara's (1972) finding that reverberatory activity was indeed established in the subcerebellar loop when picrotoxin abolished the Purkinje inhibition from the cerebellar cortex. It would be interesting to conduct an analogous experiment by blocking inhibitory transmitters in the tectum.

The type of model presented in this section is generic, and must be specialized in different ways in the light of differing ethological, anatomical, and physiological observations on different species. Moreover, the experimental literature at any time is not only not complete, but is also not free from disagreements. The job of the modeler, then, cannot be to provide a model constrained entirely by the current set of parameters, nor can or should modeling wait till "all the parameters are in." Rather, our task is to work closely with experimentalists in providing models that better structure the data, so that we can assess which families of data can be integrated in a consistent fashion. With respect to the issues raised in this section, further developments will require more refined analysis of the retina and of the temporal response of the different types of tectal neurons. Other contributions to the further "evolution" of *Rana computatrix* involve studies of predator-avoidance and stimulus-specific habituation, as well as the work to be discussed in the next section. Both experimental and simulation research are necessary and will clarify the real nature of the neural mechanisms of visuomotor coordination.

5. Layered models for depth perception

We now look in more detail at the schemas for depth perception developed by House to provide the subsystems for the "worm depth map" and "barrier depth map" for the detour model of Arbib and House (1987) discussed briefly in Section 3. To provide background for out toad models, we first consider the problem of stereopsis, or segmentation on the basis of depth cues, in humans and monkeys. To test the hypothesis that depth computation precedes recognition, Julesz (1971) developed a method based on random dot stereograms. These are binocular stimulus pairs each of which contain only visual noise but so designed that the patches of visual noise presented to one retina are identical to, but at varying disparities from, patches presented to the other retina. Julesz found that human subjects' visual systems were able to carry out the appropriate binocular matching to perceive displaced surfaces, stippled with random patterning, at varying depths in space. In other words, without precluding that some depth perception could follow pattern recognition, Julesz did establish that the formation of a depth map of space could precede the recognition of pattern. He offered a model of this process in terms of cooperative computation involving an array of magnetic dipoles connected by springs.

For the brain theorist this raises the question whether the depth map could be computed by a cooperative process involving realistic neurons (assuming that the data on what constituted realistic neurons was provided by the work of Barlow et al. 1967). One of the first papers to address this issue was that by Arbib et al. (1974) (Figure 21), who built a cooperative computational neural net model for constructing the depth map "guided by the plausible hypothesis that our visual world is made up of relatively few connected regions." The neural manifold of this model had cells whose firing level represented a degree of confidence that a point was located at a corresponding position in three-dimensional space. The neurons were connected via inhibitory interneurons to embody the principle that cells that coded for nearby direction in space and similar depth should excite each other, whereas cells that corresponded to nearby direction in space and dissimilar depth should inhibit each other (see Sperling [1970] and Nelson [1975] for related models). It was shown through computer simulation by Dev (1975) and later established by mathematical analysis by Amari and Arbib (1977) that this system did indeed yield a segmentation of the visual input into connected regions. Later, a variant of this model was published by Marr and Poggio (1977), and in subsequent writings Marr showed how our plausible hypothesis that the visual world is made up of relatively few connected regions

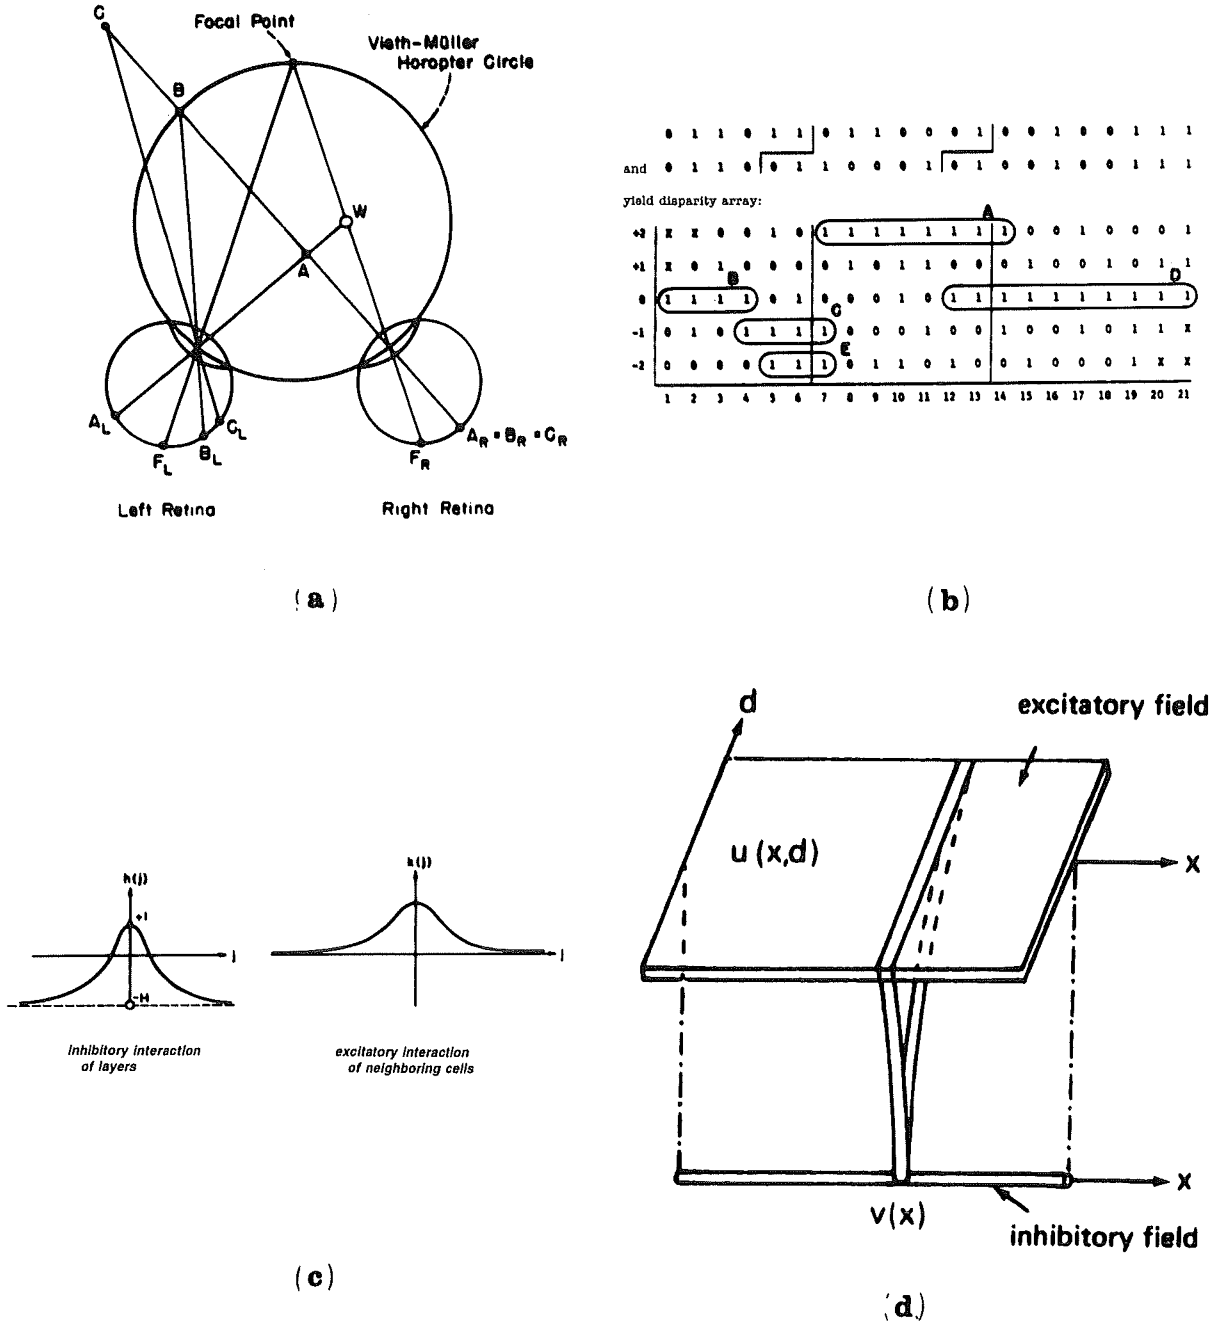

Figure 21. (a) Points projecting to the same point of one retina are projected to points with different disparities on the other retina. (b) The problem of resolving ambiguity: We conceptualize "layers" of cells (they are really in "columns"), one for each gross disparity. The aim is to segment the activity into connected regions. (c) Coupling coefficients for one approach to the problem: moderate local cross-excitation within layers; increasing inhibition between layers as the difference in disparity increases (from Arbib et al. 1974). (d) The full model of competition and cooperation, which allows the idea shown in (c) to be subject to mathematical analysis (from Amari & Arbib 1977).

could be developed into an elegant mathematical theorem relating the structure of a depth-perception algorithm to the nature of surfaces in the physical world.

With this work, then, it was established that depth maps could be constructed by a method of computation guided by the hypothesis that the world was made up of surfaces, and that the algorithm could involve some form of cooperative computation. The cooperative computation algorithms discussed above exhibited the problem of false minima, however. For example, consider a picket fence. Suppose that, by chance, the initial activity favors that match of fenceposts i, i + 1, . . . , i + j as seen by the left eye with fenceposts i + 1, i + 2, . . . , i + j + 1, respectively, as seen by the right eye. In the cooperative computational model, this initial mismatch could co-opt the possible choices of neighbors and end up with a high-confidence estimate that the fence was at a different depth from that at which it actually occurred. This provides a local "energy minimum" for the algorithm. The question then arises of how one could come up with an algorithm that would avoid at least some of these false minima. The answer provided by Marr and Poggio (1979) can be related to two contributions, one from machine vision, and one from psychophysics: the idea of pyramids or processing cones, and the idea of spatial frequency channels, respectively. Marr and Poggio developed a system in which, with hardly any cooperative computation, a fairly confident rough-depth estimate for different surfaces could be made using the low spatial-frequency channel. This rough model was then used to control vergence eye movements, sculpting a more detailed spatial map on the first approximation through the disparity information provided via channels of higher spatial frequency. In mammals, the disparity information actually seems to be precise only over a very small depth range. Disparity cues must be augmented by vergence information. Work in both human psychophysics and monkey neurophysiology (Poggio & Poggio [1984] review these and other approaches to stereopsis) has suggested that primates use three broad classes of depth detectors: those tuned for fine-depth disparities, those tuned for stimuli nearer than the coarse fixation distance (near cells), and those for farther stimuli (far cells). The hypothesis is that near and far neurons help initiate vergence, whereas tuned excitatory neurons guide completion and maintenance of vergence.

On the basis of psychophysical evidence that figural neighborhood interactions are involved in human stereopsis, Mayhew and Frisby (1980) conjectured that the matching processes are integrated with the construction of a primitive binocularly based description of image intensity changes. They then offered an algorithm, STEREOEDGE, whose initial stages used local piecewise binocular grouping of adjacent, similar, zero crossings or peak matches, whereas later stages used a Waltz-type relaxation process (Waltz 1975; for a review of similar models see Davis & Rosenfeld 1981). The earlier stages were much influenced by the work of Marr and Poggio, whereas the later stages exhibited the cooperative processing that Marr and Poggio (1979) sought to exclude. As Mayhew and Frisby (1981) note, both Marr–Poggio algorithms seek to select matches according to depth continuity rather than figural continuity, thus differing considerably from STEREOEDGE.

Prazdny (1987) offers an explicit critique of the Marr and Poggio (1979) proposal that false matches may be avoided by trading off resolution for disparity range using a coarse-to-fine matching strategy. He notes that low and high spatial frequencies are often informationally orthogonal. For example, if grass is viewed through a picket fence, there is no reason why the disparities of the fence should be related in any way to the disparities of the grass surface. He offers a specific laboratory test using a random-dot stereogram in which the background plane is transparent, and two depths, one from low and one from high spatial frequencies, can be observed simultaneously. He concludes that patches of the visual field may be fused and then held "locked" by some form of hysteresis as proposed by Julesz (1971). Prazdny (1985) offers an algorithm that successfully detects disparities generated by opaque as well as transparent surfaces. An interesting feature of the model is the absence of the explicit inhibitory connections employed by Sperling (1970), Dev (1975), and Nelson (1975). The principal disambiguation mechanism is facilitation due to disparity similarity; Prazdny argues that dissimilar disparities should not inhibit each other because, when there are transparent surfaces, a disparity may be surrounded by a set of features corresponding to other surfaces.

After this brief excursion into the analysis of mammalian systems, we now return to the depth schemas required for our analysis of detour behavior in the frog and toad. It must be noted that even if a specific brain mechanism is used by the brain of one animal this need not be the mechanism used by a distinct species. It is known that the frog and toad can snap with moderate accuracy at prey located in the monocular visual field; this led Ingle (1976) to hypothesize that for the frog it was accommodation (focal length information for the lens) that subserved depth perception. More detailed experiments by Collett (1977), which involved placing specially constructed spectacles on the nose of the toad, showed that the story was more complex. The monocular animal did indeed use accommodation as the depth cue, but for an animal with prey in the binocular field, the major depth cue was disparity, with accommodation cues exerting perhaps a 6% bias on the depth judgment based on disparity. The problem, then, is to design a model that will function on accommodation cues in the monocular animal but will nonetheless be most dependent upon binocular cues within the binocular field.

With this, we turn to two models of depth perception in the toad due to House (1982; 1984). First a note on coordinate systems. In describing the disposition of prey and barriers, it is customary to use a cartesian grid fixed on the floor of the animal's arena. In describing the visual input, we project the three-dimensional world radially onto the retinas, thus providing the animal with *two* different two-dimensional maps of the world. As is clear from Figure 6A, a given radial direction mapped on one eye may signal a turn to the left for a nearby target, a turn to the right for a farther target. We thus posit that the toad's depth map *encodes* the position of external stimuli in a body-centered set of radial coordinates, one for the turn angle required to fixate an object and the other for depth, with acuity decreasing with distance from the animal.[1] However, the nature of that encoding in neural tissue is still unknown. In House's first model, which is our current model for depth mapping of barriers, there is an explicit code, with a small patch in the neural map encoding a single angular direction but a full range of depth and the proportion of cells representing nearer depths much greater than the proportion of cells representing farther depths. In House's second model, our current model of prey localization, there is no explicit map. Following Collett and Udin (1983), it is posited that each tectum localizes a prey target on its two-dimensional map, and it is then up to the motor system to compute the three-dimensional location from the disparity of these two signals.

House's first model, the Cue Interaction Model (House 1982) uses two systems, each based on Dev's (1975) stereopsis model, to build a depth map. One is driven by disparity cues, the other by accommodation cues, but corresponding points in the two maps have excitatory cross-coupling. The model is so tuned that binocular depth cues predominate where available, but monocular accommodative cues remain sufficient to determine depth in the absence of binocular cues. The model produces a complete depth map of the visual field, and so is

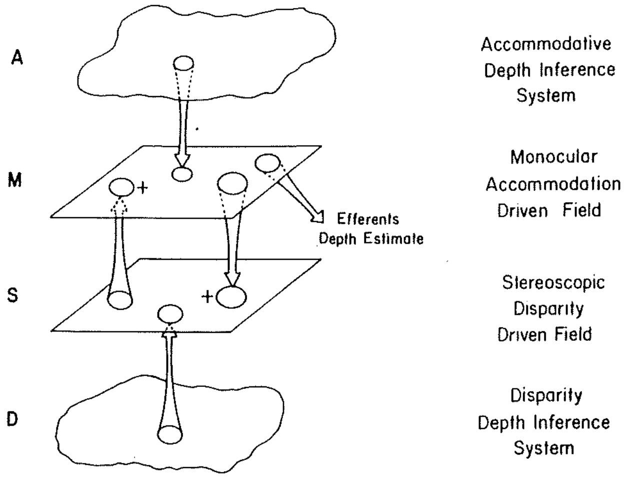

Figure 22. A depth-mapping model using cross-coupling between an accommodation-driven system and a disparity-driven system (House 1982).

appropriate for building a representation of barriers for use in navigation. The model is shown in Figure 22. At the top we see an accommodation-driven field, M, which receives information about accommodation and which – left to its own devices – would sharpen up that information to yield relatively refined depth estimates. Below, we see the type of system, S, posited by Dev to use disparity information and suppress ghost targets. However, the systems are so intercoupled that a point in the accommodation field M will excite the corresponding point in the disparity field S, and vice versa. Thus a high confidence in a particular (direction, depth) coordinate in one layer will bias activity in the other layer accordingly. The result is that the system will converge to a state affected by both types of information – although the monocular system can, by itself, yield depth estimates.

Figure 23 shows stages in the processing by this model of a scene comprising a fence and two worms. The left-hand column of each subfigure shows the accommodation (above) and the disparity field (below) for the fence information; the right column shows prey information. In the top image of Figure 23b, we see the initial state of the accommodation field. The information is blurred, representing the lack of fine-tuning offered by accommodation. Below, we see the initial state of the stereopsis field. The targets are better tuned, but they offer ghost images in addition to the correct images. Figure 23f shows the outcome of the interaction posited in the model. We see that virtually all the ghost fence targets have also been suppressed. In addition, we see that the accommodation information has been sharpened considerably. The information is now precise and unambiguous, and thus can be used to guide the further behavior of the animal.

The above model is of interest in its own right – as a model specific to the study of the amphibian (we postulate that such a system processes barrier information), and as an indication of the class of stereopsis models based on multiple cues. Here, however, I would like to stress its more general significance. We claim that *cooperative computation* is a general principle of brain operation, with different sensory systems providing different patterns of information to be factored into the determination of the overall behavior of the organism. Since we explicitly designed the model of Figure 22, we know that one layer represents accommodation information whereas another represents disparity information – and we can clearly see the differences in these types of representation in Figure 23b. However, Figure 23f represents the sort of state of activity that is much more likely to be seen during the ongoing behavior of the system, and here we see that both surfaces represent pooled information based on the interaction between the layers, rather than information directly supplied by sensory systems. *This clearly indicates the dangers of experimentation based on feature analysis without related high-level modeling.* As we can see, feature analysis of Figure 23f would simply show cells responsive to information available at a specific depth and visual direction. Only a subtler analysis, guided by a model of the kind presented here, would allow the experimenter to discover that, although much of the time the two surfaces exhibited congruent activity, one was in fact driven primarily by accommodation, while the other was driven primarily by disparity information.

Collett and Udin (1983) showed that, for the task of unobstructed prey catching, toads are able to make accurate binocularly based depth estimates even after n. isthmi (NI), the major cross-tectal binocular relay, has been lesioned. Collett et al. (in press) report behavioral studies on two toads with lesions that destroyed or disconnected most of both NI. To find out whether binocular cues remain effective after NI lesions, they tested prey-catching behavior when the toads viewed prey through prisms that changed horizontal binocular disparities or through convex lenses that altered the accommodative state of the eyes. In both cases, there is a conflict between monocular and binocular cues. As we have seen, Collett (1977) found that binocular cues predominate in the normal animal; the present study showed this also to be true in the NI-lesioned toads. Collett and Udin postulated that the toad may use triangulation to locate the prey, rather than a process of disparity matching, much as the mantid has been hypothesized to form depth estimates by comparing output signals from the two optic lobes (Rössel 1983). They also found that toads undershoot their prey equally whether the disparities imposed by prisms are horizontal, vertical, or oblique – contrary to mammalian disparity detectors, which can operate only if there is reasonable vertical alignment between stimuli on the two retinas. In particular, Collett et al. doubt that a point in the visual field is resolved by the tectum into its horizontal and vertical components. Rather, they offer for consideration the notion that the tectum codes position in polar coordinates, with disparities measured as the difference between the radial coordinates of a point in each eye.

House's (1984) second model, the Prey Localization Model, incorporates the triangulation hypothesis. Each side of the brain selects a prey target based on output of the contralateral retina and computes a depth estimate by triangulation to adjust lens focus. If the selected retinal points correspond to the same prey object, then the depth estimate will be accurate and the object will be brought into clearer focus, "locking on" to the target. If the points do not correspond, the resulting lens adjustment will tend to bring one of the external objects into clearer focus, and the two halves of the brain will tend to

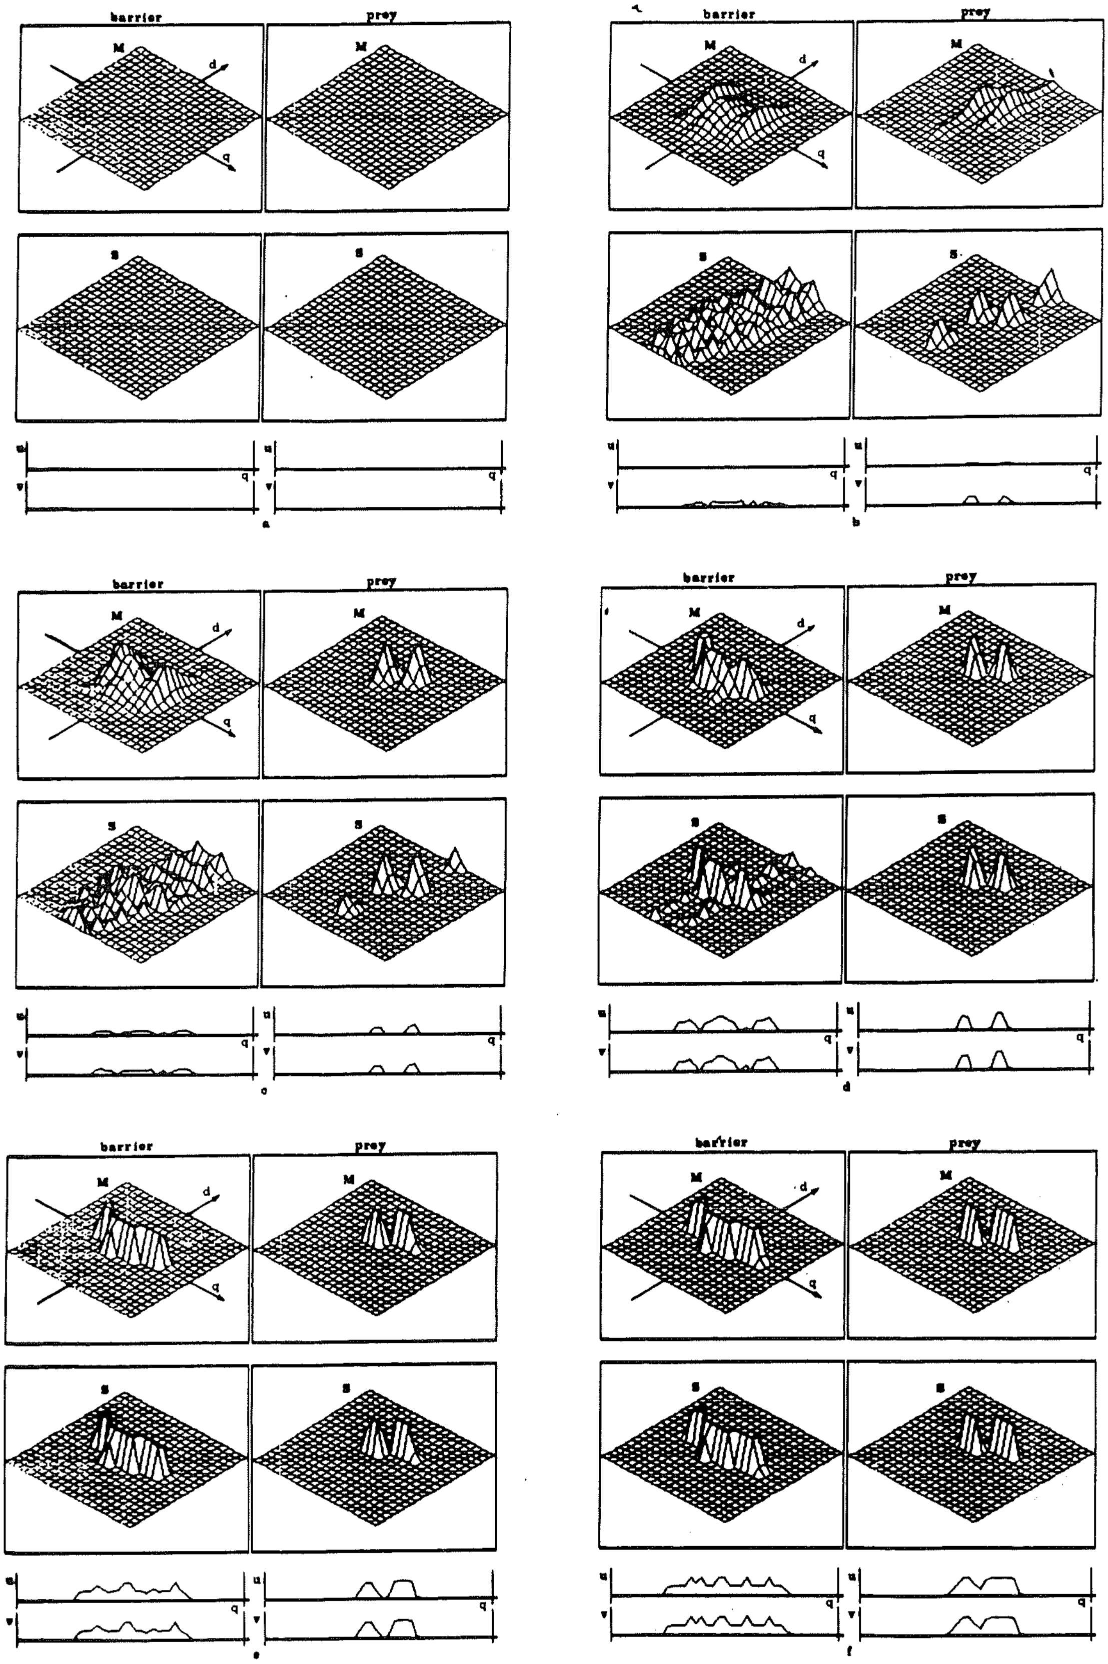

Figure 23. Time course of the model: The time course of the depth model from its initially inert state (a) to a satisfactory depth segmentation (f) is shown here. All figures are in the retinal angle versus disparity coordinate system. Successive figures are temporally spaced 1.4 field time-consonants apart. Thus, the elapsed simulation time represented is nearly 7 time constants. The two-dimensional grids show the level of excitation of the various fields, and the line graphs under the grids indicate the intensity of localization on the retinal angle axis of excitation in the inhibitory pools (House 1982).

choose that object over the other. However, Caine and Gruberg (1985) find that frogs with lesions of NI failed to respond to either threat or prey stimuli in the corresponding region of the visual field (contradicting Collett and Udin), while exhibiting normal barrier avoidance and optokinetic nystagmus.

All this poses a clear challenge for future theory and experiment.

6. Conclusions and implications

I have argued that computational models of cognition should yield a functional model of how behavior is achieved through the interaction of a number of simultaneously active computing agents called *schema instantiations.* The dissection of a behavior in terms of interacting schemas is to provide the meeting ground for the behavioral and brain sciences. For the cognitive scientist, a schema-theoretic model may suffice so long as it survives the tests of the psychological laboratory. For the neuroscientist, such a model must be refined to hypotheses about the localization of schemas, constrained by the study of brain lesions and of detailed neural circuitry.

One useful level of analysis intermediate between the purely functional schema and the individual neuron is that of the layer in which we can represent activity as a spatiotemporal function without quantization in terms of the activity of separable neurons. Yet it can be valuable to model schemas, their instantiations and their interactions, by interacting layers of neuronlike elements or by nets of "intermediate-level" units, even if the nets are little constrained by anatomy or physiology. They still serve to extend our vocabulary for discussing the properties of neural networks, better preparing us to handle new data as they become available.

Our models stress *cooperative computation* of interacting subsystems at two levels of analysis: (a) competition and cooperation in closely coupled layers of neurons and (b) the cooperation of different schemas, each with partial sources of information, in committing the organism to a course of action.

Our focus on visually guided behavior in the frog and toad provided core examples of schema models, of models based on interacting layers of "neuronlike" units, and of neural network models closely coupled to detailed data from neuroanatomy and neurophysiology. Besides illustrating these types of models and the give and take between them, the examples were also meant to suggest the excitement of incrementally "evolving" an integrated account of a single animal, seeing the challenges posed by combining different aspects of vision with mechanisms for the control of an expanding repertoire of behavior. I trust it is clear that the results have general implications for the study of the neural bases of behavior that are in no way limited to the frog and toad. In Section 5, the "disparity" between the findings of Caine and Gruberg (1985) and Collett and Udin (1983) on the effects of nucleus isthmi lesions on depth perception and prey catching highlighted the need for the *incremental style of modeling,* which allows us to gain insight into current experimental data yet leaves us the flexibility to learn from new data, to make sense of conflicting data, and to offer both new hypotheses and new challenges for experimental testing.

The last point may be emphasized by drawing an analogy between *Aplysia* and *Rana computatrix.* It is not many years ago that the study of invertebrates was regarded as being peripheral, at best, by the majority of neuroscientists. Yet the increasing attention to mechanisms of neural function has made *Aplysia* and its squishy or crunchy cousins invaluable in the study of basic cellular mechanisms of facilitation and habituation, and in the study of the coupling of neurons for rhythm generation, for example. However, these very successes at or near the cellular level should not blind us to the drastic differences in organizational principle that separate the primate brain from the *Aplysia* nerve net. My suggestion is that if creatures that have evolved by chance can provide insight into cellular mechanisms, then "creatures" that "evolve in the computer" can provide opportunities for understanding organizational principles. Moreover, I argue that those organizational principles are not to be sought solely in terms of cellular mechanisms but in terms of structural constructs (layers and modules), functional constructs (schemas), and computational strategies (cooperative computation, competition and cooperation in neural nets, etc.). *Rana computatrix* is thus a test bed not only for the incorporation of specific data on neural circuitry but also for the development of organizational principles. Data on the frog and toad thus do not exhaust the implications of *Rana computatrix.* Rather, the better we understand the relation of detailed neural circuitry to models that are more schematic (in both senses of the word), the better we can adapt these models to provide insight into analogous systems in other organisms.

To close this discussion, let me note that the analysis of levels here has implications for the social structure of cognitive science in general and brain theory in particular. Models evolve in complexity in (at least) two ways:

(a) through replacing units by more detailed networks; and

(b) through integrating more and more schemas into explanations of an increased range of behaviors and/or brain regions.

A problem with (a) is that the simulation of a large array of units may not be feasible if each unit is described in the most refined way currently available. A problem with (b) is that the analysis of one "macroschema" may itself be a full research task, yet the interfaces with other schemas may introduce constraints that simply cannot be ignored. The resolution will require much greater "cooperative computation" between research groups, and the development of standards for interfacing their work. Here are three "scenarios":

1. We discussed the "evolution" of two different, but overlapping, subsystems of *Rana computatrix:* one for prey selection, and one for depth and detours. We must now synthesize these into an overall system. In the process, we may better understand some of the complexity of tectal circuitry as we probe its computations in subserving several rather different schemas.

2. Much of our modeling addresses the ways in which retinotopically structured subsystems may interact to determine the location of objects or task-related features to commit the organism to some spatially structured course of action. In the meantime, Robinson and his colleagues (e.g., Robinson 1981) have used control theory

to explicate how the (numerically, rather than spatially, encoded) position of a target can be transformed to commands for muscle contractions to bring the eye to bear on the target. This raises fascinating problems in the interfacing of such efforts:

(a) What control system for, say, snapping at prey, would be analogous to Robinson's systems for eye movements?

(b) What constraints on such a control system are imposed by making it accept retinotopic, rather than numerical, inputs?

(c) How can a better appreciation of variables involved in muscle control change our conception of the proper output for the "more or less visual" portions of the visuomotor control system?

To the extent that different groups work on these problems, it would be most helpful to have an agreed-on "interface language" to express the flow of information between subsystems. This should be so designed that researchers could couple programs developed in different laboratories and use the experience to suggest changed input/output specifications to constrain model development at each site.

3. In our versions of *Rana computatrix* to date, we have used "black-box" models of the retina, but we now seek more insight into the details of spatiotemporal patterns of retinal output (Lee 1986). To this end, we are simulating a 10,000 neuron network on a CRAY supercomputer. Meanwhile, other workers find it necessary to exploit the full power of a VAX 11/780 to explore the detailed synaptic and membrane interactions within a single neuron. Here, the suitable interface seems to be a description of the cell's input/output relations at an intermediate level of complexity. At any time, the task of one group would be to test how well a network of such model neurons would explain overall retinal response, while another group would explore how adequately such a representation corresponds to the fine details of neuronal form and function.

As a corollary to these methods should come far better tools for the exchange of data between experimentalist and theorist. Published data often lack details crucial for model testing and development. In the future, we should see increased availability of large data sets recorded in forms that will allow easy processing for detailed comparison of model and experiment. More generally, the self-conscious analysis both of what levels of analysis are appropriate for cognitive science and of explicit methodologies for interfacing between them should greatly aid the evolution of "cognitive science" from its current status as a loose federation of disparate subjects into a genuinely integrated field. My point is that the correct strategy is not the reduction of cognitive science to one "fundamental level" but rather the development of strategies for "two-way reduction" (Arbib 1985; Arbib & Hesse 1986) whereby the reality of different levels may be recognized and reconciled to create a richer whole.

APPENDIX A

Perceptual schemas for high-level vision. Appendix A outlines the approach to schemas for high-level vision adopted by Weymouth (1986) as part of the VISIONS project on machine vision (Hanson & Riseman 1978). The system is to automatically go from a color photograph (e.g., 256×256 pixels with 8 bits per color) to an interpretation network describing the scene, with pointers from object-nodes to the corresponding region of the image.

Table 1. *Object classes for interpretation*

- sky
- foliage
- tree crown
  - bush
- ground plane
  - grass (lawn)
- other ground cover
  - road
    - paved (called "dark" road)
    - dirt (called "light" road)
- buildings
  - other buildings
  - house
    - roof
    - walls
      - wall features (or "markings")
        - window group
          - window
          - shutter
        - door

Low-level processes work independently of object-specific knowledge to extract regions of the image to serve as candidates for surfaces of objects, to extract long lines, and to measure features – such as color, texture, shape, size, and location – of these regions and lines. (Other processes can yield further information, such as depth and motion, and the present methodology can make use of these powerful cues.) The result is an intermediate representation of the image that can serve as the basis for interpretation processes that apply object-dependent knowledge to label regions with interpretations, grouping regions and building descriptive networks in the process. Table 1 shows the hierarchy of objects for which knowledge is encoded in the present system.

The result of interpretation is a hierarchical network comprising a segmentation of the image into interpretable regions, a labeling of the regions, a geometric model of the scene describing object placement and structure, and a network containing the spatial, compositional, and taxonomic relations of objects in the scene and in its image (Figure A1.1).

The actual system not only produces the network but also a confidence value associated with each node. Interpretation integrates many procedures: matching 3D models of geometric structure to line segments in the image; using pattern identification techniques to identify classes of objects associated with regions; using a network of object-part relations to guide the process. The system uses parallel distributed control, taking advantage of redundancies to recover object identity from noisy error-ridden data.

It will help to distinguish four types of entity (Figure A1.2): A *schema* encodes knowledge (both local and contextual) of criteria for recognizing an object of a particular kind. A *schema instance* is the trace of running a schema on a particular subset of (the network constructed on the basis of) the image. When a schema instance is activated, it is with an associated area of the image and an associated set of local variables. Multiple instances would normally be associated with separate portions of the image. A *hypothesis* is the assertion that a particular object provides the interpretation of a portion of the image. It is registered when the relevant schema instance achieves a threshold in its confidence level, and will then include parameters

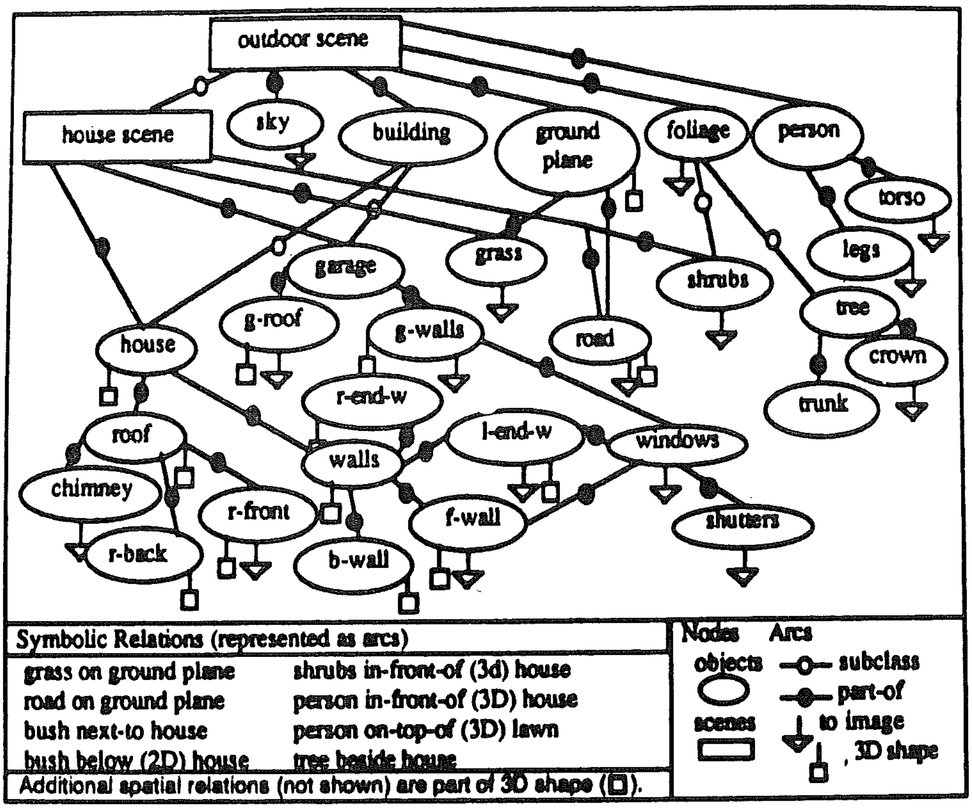

Figure A1.1. A hand-generated interpretation network for an outdoor scene containing a house, a person, a road, foliage, grass, and sky. The arcs representing the symbolic relations listed at bottom left of the figure are omitted for clarity. Note that objects may be rooted in other objects, in actual regions of the image, or may be confidently inferred (such as the back wall of the house) from the knowledge of 3D shape and structure of objects accessed by the interpretation process (Weymouth 1986).

descriptive of the object so recognized – contrast schema activation (when an instance starts to process) with schema "firing" or propagation (when an instance posts a hypothesis that can affect the activity of other schemas). A schema instance may set as a *goal* the confirmation that a certain context applies. Posting a goal may in turn lead to the forming of a schema instance to check whether a posited object occurs in a certain portion of the image.

As the process of interpretation proceeds, it may select low-level processes to, for example, refine the segmentation of some regions, or merge others. The knowledge required for interpretation is stored in what is called LTM (long-term memory), in the form of a network of schemas – as distinct from STM (short-term memory), the network that describes the particular scene. It echoes the structure of LTM, both that between schemas (for interobject and interpart relations) and that within a schema (for geometric relations of parts).

The STM network makes context explicit: Each object represents a context for further processing. The control mechanism must be able to handle multiple instances of an object class, keep track of several objects, and use to advantage the relations among objects. Interobject relations include the part–whole relation and class membership. Commonalities between two schemas (e.g., "house" and "barn") can be handled by making them subordinate to a more general schema (e.g., "building"). Other relations (all expressed as edges in the networks) concern proximity and coexistence: e.g., "attached to," "near," "must occur with," and "can occur with." Viewpoint-specific relations include "above," "below," "to the left (right) of." Scene schemas are "framelike" in that they describe the types of objects that may occur in a scene of the given type, together with their likely spatial relationships. Object schemas may either be primitive, in that they define criteria for object recognition directly in terms of properties of the intermediate representation; or composite, in which case they define the object in terms of relationships between constituent parts.

Each schema combines declarative knowledge (e.g., description of the 3D geometry of an object, or typical appearances of object surfaces, or object-subpart relationships) with procedural knowledge in the form of object-recognition routines and control programs. Object descriptions relate image events to object labels. They are available at several levels of resolution, and through a hierarchy of subparts. Hierarchies of specialization (class-subclass) and composition (part-subobject) are used as a framework for control. Much of this knowledge about an object is represented in interpretation strategies embedded within the correponding schema. Under the guidance of the schema network in LTM, and on the basis of the intermediate representation (which may itself change in response to requests to low-level processes), the interpretation processes build an interpretation network that comprises schema instances (what we have called instantiations), goals, and hypotheses linked to each other and to groupings of image features and the 3D structures that interpret them (Figure A1.3). There may, of course, be several instances of an object-class in a scene. As nodes are created in STM, their subordinate part (and class) nodes are checked for consistency with the network in LTM. Missing, but required, subparts detract from the confidence score associated with the object. Interpretations that overlap in the image that are not related by subpart or subclass arcs are discounted (mutually) because they are assumed to be in conflict. This is handled through a scanning mechanism.

In the present system data are treated as primitive hypotheses; hypotheses are made and confirmed by the selective

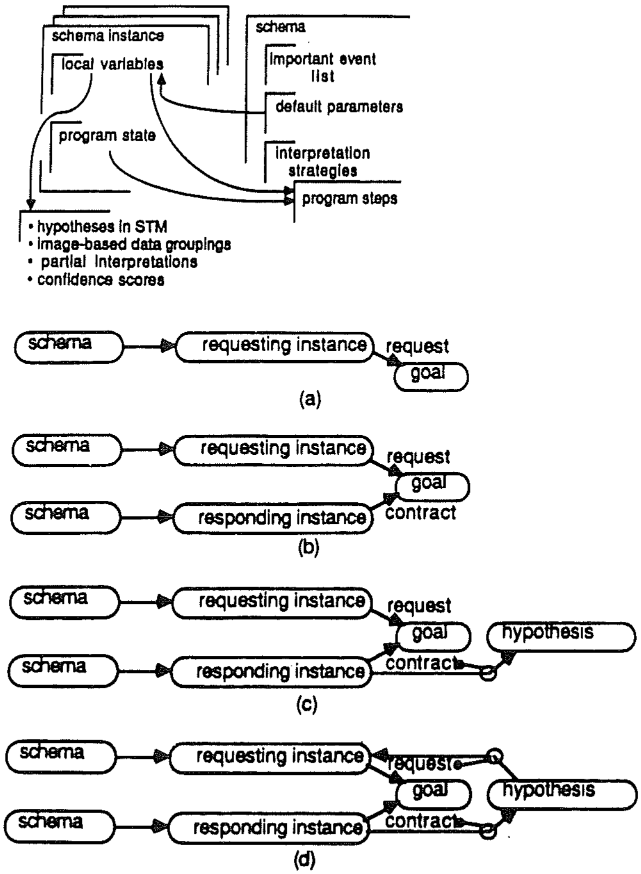

Figure A1.2. (Top) A schema instance in STM (short-term memory) includes the complete program state of one invocation of an interpretation strategy (which need not be serial) from a schema in LTM (long-term memory), providing both the values of local variables and a pointer to program steps in the interpretation strategy. When the schema is instantiated, some local variables may be preset to initial default values supplied by the schema, whereas if a goal activates the schema, initial values may be set from the goal. Partial interpretations are built up by pointers to hypotheses in STM. (Bottom) The relationship between schemas, instances, goals, and hypotheses. When (a) a schema instance requests a goal, a request link records that relationship. The goal then activates a new instance of an appropriate schema (b), with a contract link showing which goal the instance is working on. If a hypothesis to satisfy the goal is created by that instance (c), this hypothesis is attached to the contract link and posted in STM. A flag on the contract link signals whether further processing is possible. The hypothesis is then posted (d) to the instance on the request link, which may resume any processing suspended pending satisfaction of the requested goal. Unless this instance terminates the connections, the various links may remain in place as pathways for further communication (Weymouth 1986).

organization of prior hypotheses; processing of a hypothesis is determined by a rating estimating its effectiveness in, and consistency with, the interpretation process. Each schema refers to an object. Not only can a schema create the hypothesis for the object, it can also reorganize related hypotheses in the context of the object or scene represented by the schema. Activation of a schema may ultimately cause the creation of several hypotheses, either the alternate organizations of related hypotheses or multiple instances of objects. The schema includes processes that encode the context represented by the schema in the control of perception.

A schema contains a list of important events, a set of programs, and a set of parameter values. When a schema is instantiated, a specific set of the schema's programs (which are run in parallel) are invoked. Default settings of parameters may be overridden by parameters supplied when the instance is invoked. The programs, executed in parallel with equal priority, control the invocation of functions to extract data and data structures, and to create and manipulate hypotheses. For example, the roof schema invokes a function called roof-rectangle-builder that tries to combine image regions to build a (perspective view of) a rectangle within appropriate size, color, and texture constraints. This function invokes more primitive functions for the organization of data, creating intermediate data structures such as corners and parallel lines. Instantiation of a schema may be either data-driven or goal-driven.

In data-driven (bottom-up) activation, any event from the "list of important events" can activate the schema. To initiate processing, each region in the intermediate representation has an associated list of adjacent regions and lines; each line has a list of regions it borders. Each line and region has a unique index. We distinguish lines and regions as properties extracted from the image, from edges and surfaces as items in the scene. Lines and regions can be filtered bottom-up for significance, so that certain ones will then have edge and surface hypotheses, respectively, associated with them. These provide "islands of reliability" to start up the schema activation process. Furthermore, schemas can associate surfaces and edges with objects, thus elevating certain regions and edges to become primitive hypotheses in the network describing the scene.

An example of a data-driven interpretation routine for the sky is as follows. (Other routines can use contextual cues to determine that a region is skylike.) Five measures of goodness of fit are calculated for a region:

$m_1(r)$ measures the extent to which the location is middle to high in the image;

$m_2(r)$ measures the extent to which the color of r is high intensity, high saturation, and high blue (more generally, it might also measure closeness to "cloud color" or "smog color");

$m_3(r)$ measures the extent to which r is a large, wide region;

$m_4(r)$ measures how far the texture exhibits low blue deviation;

$m_5(r)$ evaluates the shape of r in terms of an optimal height-to-width ratio near one-half.

Then a region is adjudged more "skylike" the greater is

$$m_{sky}(r) = 5.5m_1(r) + 4m_2(r) + m_3(r) + m_4(r) + m_5(r).$$

The particular m_i's and the weights were chosen by trial and error – the larger the weight, the more crucial the feature. Future work will show how to automate their extraction from a large collection of hand-scored images. The point here is not to advocate the particular measure chosen, but simply to note that such measures can be effective in assigning a confidence level to an initial "data-driven" classification of a region on the basis of local cues. In general, this will only be the first step and in no way forces the final interpretation. As we see from this example, some object classes have known limits on surface color and texture; in other cases, geometric form and surface appearance may be more useful.

For another example, a long horizontal line below the sky could activate the roof schema. With additional evidence, the roof hypothesis might be formed. This might be on the important event list of the house schema, thus instantiating it. A chain of bottom-up activation leads to an early interpretation, which may then provide a context for the goal-driven extraction and interpretation of detail.

In goal-driven (top-down) instantiation, a goal is created by some active program, and supplied with certain parameters. A schema that can potentially "satisfy" that goal gets activated and "attaches" a schema instance to it. This instance runs its programs. If a level of sufficient confidence is reached, "satisfying"

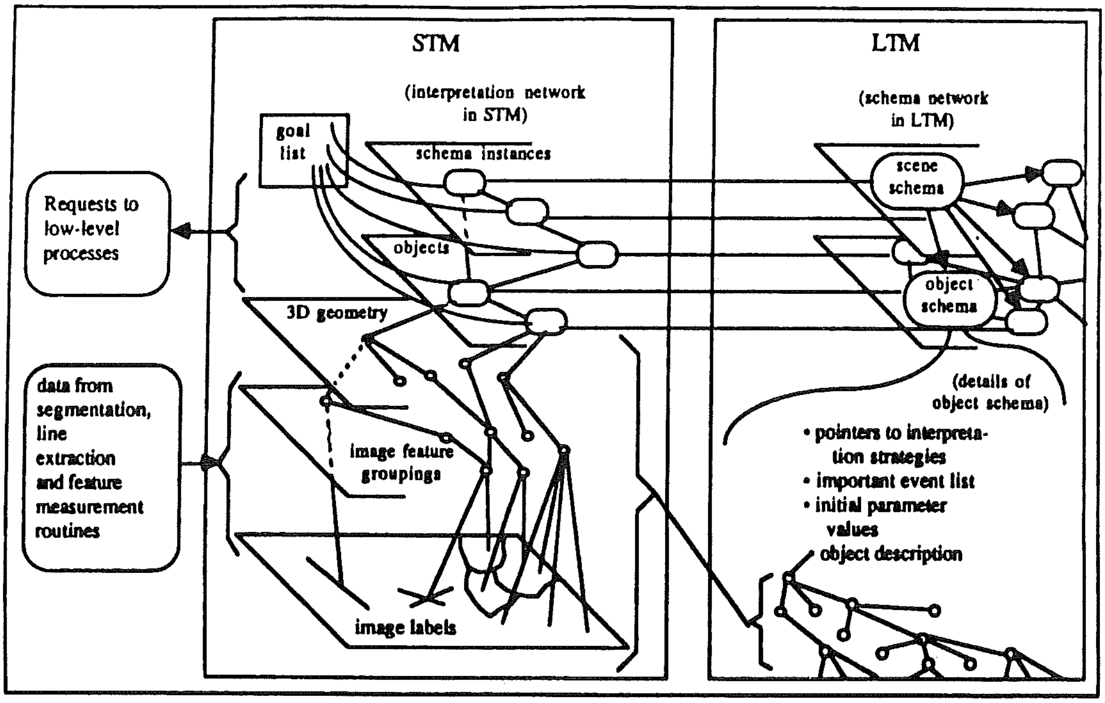

Figure A1.3. Interpretation strategies are stored in schemas that are linked in a schema network in long-term memory (LTM). Under the guidance of these schemas, the intermediate representation is modified and interpreted by a network of schema instances that label regions of the image and link them to a 3D geometry in short-term memory (STM) (Weymouth 1986).

the goal, a hypothesis is created, and a message is passed to the goal to tell it that it has been satisfied. The instance may then stop processing (waiting for further instructions) or may continue, attempting to refine or change the hypothesis. A cycle of incremental improvement can continue until (if ever) the schema instance has reached the end of all its programs.

When a schema program S executes a "wait-for-hypothesis" statement, it suspends execution until a hypothesis of that type is created. Alternatively, if the requisite hypothesis has an associated schema, then S can instantiate the associated schema, specifying the image area and initial values for the local variables. The new instance treats the request for activation as a goal for a hypothesis of the object represented by the schema (Figure A1.2). If the hypothesis is created, the schema instance will signal the satisfaction of the goal. The activating schema may either continue processing or suspend activity pending the receipt of this signal. In this way schema instances can be coordinated to create structures within which hypotheses are linked (through relations) in the final interpretation. Each schema unifies multiple strategies for the recognition of a specific object class; these strategies serve the same goal and interact strongly.

Early matches are used to construct a skeletal description, which can in turn be used to guide the search for further facts. The facts that are easier to match, and how they can be used to limit search, have to be encoded as domain-specific knowledge. Just which cues are most potent will depend on the specific image. Context provided by a few quickly perceived cues leads to the activation of a schema assemblage, which in turn guides future applications of perceptual schemas. For example, a roof might be recognized primarily by its geometry if it is unoccluded but by its spatial relation to an already recognized wall if it is behind some trees.

Procedural information includes the order in which object primitives are best perceived (edge between roof and sky matches image better than edge between roof and wall), and the relative perceptual importance of subparts (look for the windows before the doors). The object description provides structures for integrating representations of geometric structure and image appearances: a compromise between a complete geometric description (possibly with free variables) in a 3-D object-centered coordinate system and a summary of expected appearances. For example, with a tree crown or crumpled paper, only the latter may be useful. The description of object appearance is to be related to feature values (such as color, texture, and shape) of regions in the segmentation.

Important constraints used in the design of the schema system characterize when two object labels can be applied to the same space. Constraints are provided by a part–subpart hierarchy, which allows the description of complex subparts (e.g., "roof" and "house") to be expressed only once and localized, with surface and line used as primitive subparts; and there are specializations (e.g., "house" from "building") forming a class–subclass hierarchy, which allows the description of a more general prototype to be associated with the object. Each specialization of a generic description inherits properties of the generic object – unless superseded by the specification of particular features on relations.

Relations like "near," "above," "to-the-left-of" are used with respect to the frame of reference of the camera, whereas relations concerned with assemblages of objects include "attached-to" and "occurring with." For example, a gable roof may be characterized by two equal rectangles meeting at a common line (the crest of the roof), meeting at an angle between 25° and 105°, and having a certain range of possible values of texture and color. The roof description (Figure A1.4) has nodes for object, surface, edge. It is assumed that the properties of the roof side that is not visible in a particular view are the same as those of the visible roof side. Each view of the roof can be characterized by, for example, grouping of lines and the relations among them.

Interpretation must amalgamate many uncertain results into a more perfect union. During visual perception, we may find some interpretation mistaken (that's *not* a shutter) yet retain the parts of the interpretation consistent with the new data (it *is* rectangular). Many plausible interpretations proceed in parallel, in such a way that the most likely one dominates, sharing

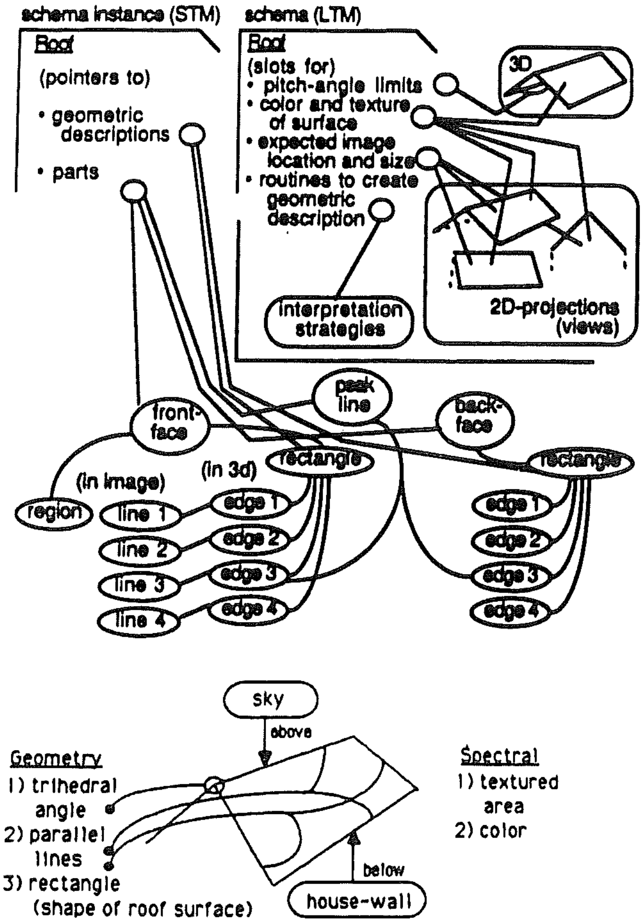

Figure A1.4. (Top) A network describing a roof: The STM structure contains both items based on image data and the derived three-dimensional structure. The nodes may have actual parametric values inferred from a geometric model. The interpretation network is derived from the schemas in LTM that contain information on both the 3D spatial relations of the object parts and on various projective views. (Bottom) Several factors contribute to the interpretation of a portion of the image as a roof including region characteristics (color and texture), shape information from line data, and relations with other objects (sky and other house parts) (Weymouth 1986).

common data structures and descriptions. The structure that is selected as the interpretation is just another organization of the facts and intermediate representations. Context switching consists of attending to a different portion of the structure that accompanies the selection of an alternative explanation as the most likely one.

An example of interpretation is given in Figure A1.5. The reader is referred to Weymouth (1986) for further examples and a much fuller specification of a range of schemas, including those which elicit further low-level processing to provide necessary refinement of the intermediate representation, and those which construct geometric models of objects in the scene.

APPENDIX B:

Motor schemas for dextrous hand control. Appendix B illustrates the notion of motor schemas by outlining the approach to schemas for control of dextrous hands adopted by Lyons and Arbib (in press). Lyons (1986) developed a generic hand simulation adapted both (1) to the human hand and (2) to the Salisbury (JPL/Stanford) robot hand, which has three fingers, each with three degrees of freedom (but four motors). The approach confronts two fundamental difficulties in dealing with complex robotic systems: the coherent control of many degrees of freedom, and the integration of sensory information into motor behavior. As in Section 2, the approach is based on perceptual schemas and motor schemas. A complex domain of interaction (tasks and/or perceptual events) is represented as a schema assemblage or, more precisely, as an assemblage of (potentially) concurrently active schema instantiations.

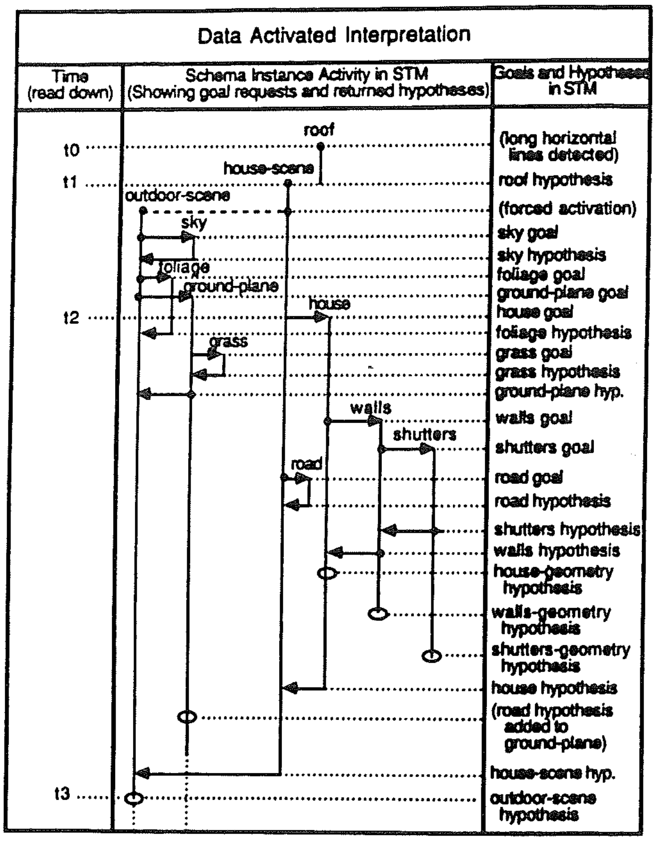

Figure A1.5. This figure shows the time relations of goal requests, schema activations, and hypothesis creations during interpretation. The vertical axis shows the passage of time, starting at the top. We see schema instances created by goal requests (except for the initial activation of the roof schema by image data). Note that much of the ordering of hypothesis formation is essentially random, the result of implementing an inherently parallel process on a serial computer (Weymouth 1986).

The semantics of a schema is a generic *port automaton* description augmented with *instantiation* and *deinstantiation* operations. The instantiation operation forms a schema instantiation (*SI*) by taking as input a schema together with instantiation parameters: (1) initial values for internal variables, and (2) a connection map specifying connections of ports on this SI to ports on other SIs. The *behavior section* of an SI cycles continuously until the SI is deinstantiated. We do not set forth the formal semantics here, but only the formal syntax, an informal indication of the semantics, and a few indicative examples.

The syntax of a schema definition is

[N (ip) (op) (v) (b)]

where

N is the identifying name for the schema;
ip is the list of input port names;
op is the list of output port names;

v is the list of internal variable names; and
b is a specification of behavior.

A behavior is given by the syntax

⟨behavior section⟩ ::= ⟨Stat⟩*

(a sequence of 0 or more elements of Stat) where

⟨Stat⟩ ::= ⟨Assign⟩| ⟨If⟩ | ⟨For⟩ | ⟨Instn⟩ | ⟨Dinstn⟩ | ⟨Forall⟩

In this list of alternatives, the "conventional" statements are

⟨Assign⟩ ::= ⟨Var⟩:=⟨Expr⟩ | ⟨OutputPort⟩:=⟨Expr⟩

which evaluates an expression based on internal variables and/or values read from an input port either to update an internal variable, or to write a value to an output port; and

⟨If⟩::= IF⟨condition⟩THEN⟨Stat⟩*ELSE⟨Stat⟩*ENDIF

which is a conventional conditional.

Instantiation and Deinstantiation are novel, and have the syntax

⟨Instn⟩::= ⟨Schemaname⟩$_{\mathrm{Inst}}${V$_o$}{C$_o$}

where

C$_o$::=⟨Couplings⟩*
⟨Couplings⟩::= Ipname←SIPName | SIPName←Opname
SIPName::=Schemaname$_{\mathrm{Inst}}$(port)

and

⟨Dinstn⟩::=STOP{Schemaname$_{\mathrm{Inst}}$}.

Instantiation creates a new instance of a schema and couples its ports to extant SIs in the specified way, whereas deinstantiation simply removes the specified instance from the SI network.

Another novel instruction is the FORALL given by

⟨Forall⟩::= FORALL Schemaname DO ⟨Instn⟩*ENDFORALL

which allows all instances of a given schema simultaneously to instantiate their own similar but private SIs.

The above definition establishes a schema in terms of a body (behavioral section). We also build up schemas recursively. An *assemblage SI* is a computing agent in which the behavior is defined in terms of the behavior of a network of communicating SIs. The syntax of an assemblage definition is

[N (ip) (op) (s)(ib)(p)(n)]

where
N is the identifying name for the assemblage;
ip is the list of input port names;
op is the list of output port names;
v is the list of internal variable names;
s is a list of component schemas;
ib lists the commands to form instantiations;
p renames ports of the component SIs as ports of the assemblage; and
n defines the port connection mappings between component SIs.

The syntax of p is given by

p::=⟨Equivalence⟩*
⟨Equivalence⟩::=⟨portname⟩=⟨portname⟩

Recalling the format [N (ip)(op) . . .] for schema syntax, we list a number of predefined schemas that form the building blocks for the robot schemas.

Position-Servo Schema

[Jmotor(desired)(actual) . . . ⟨implementation dependent⟩ . . .]

which will cause the joint to move until the actual position of the joint as reported on output port *actual* will match the desired position requested at the input port *desired*.

Tactile-Contact Schema

[Tactile()(Contact) . . . ⟨implementation dependent⟩ . . .]

which has no input port, but continually reports whether or not the touch sensor with which it is associated has registered a contact.

Separable-Environmental-Facet Schema

[SEF()(F$_1$ . . . F$_n$) . . . ⟨implementation dependent⟩ . . .]

which corresponds to a sensor or a perceptual schema monitoring some "facet" of the environment. It could thus be a sensor reporting the RBG (red, blue, green) values of a single pixel, or a house schema reporting the position, size, and other salient characteristics of a house.

Each such schema may have many instantiations in a given distributed controller for a perceptual robot.

A *task unit* is an assemblage consisting of perceptual schemas, motor schemas, and linking (coordinated control) schemas.

We use the abbreviated syntax [P-M] for a task unit comprising the perceptual schema P and the motor schema M, as in the following examples:

Joint$_i$ = [Jposition$_i$ − Jmotor$_i$]
Jointf$_i$ = [Jforce$_i$ − Jmotor$_i$]
Closejoint$_i$ = [Tactile$_i$ − Joint$_i$]

A *precondition* for a task unit is an SI, which will test for some defined conditions before making an instance of the task unit. We use the abbreviated syntax

Pre ?:[P-M]

for a task unit comprising precondition Pre, perceptual schema P, and motor schema M. For example

Mug ?:[Mugpar-Graspmug]

where Mug is a perceptual schema to recognize a mug to be picked up; Mugpar is a perceptual schema to extract relevant parameters of a mug; and Graspmug is a motor schema for grasping a mug.

To exemplify the above, we now consider (Figure A2.1) the Reach task-unit schema, with abbreviated syntax

Reach = [RObject − (MoveWrist, OrientWrist)].

First we set up the RObject schema:

[RObject (Pn On) (Pt Ot)
(handlength grotate)
(Pt:= Pn-handlength
Ot:= grotate·On)]

Handlength and *grotate* are internal variables set up on instantiation to record the length and orientation of the hand that will do the reaching. (Thus a more extended task unit could show explicitly the receipt of these variables from other SIs.) In the task unit (Figure A2.1), RObject will receive the position and orientation of the object to be grasped, through input ports Pn and On respectively, from output ports F$_1$ and F$_2$ of percep-

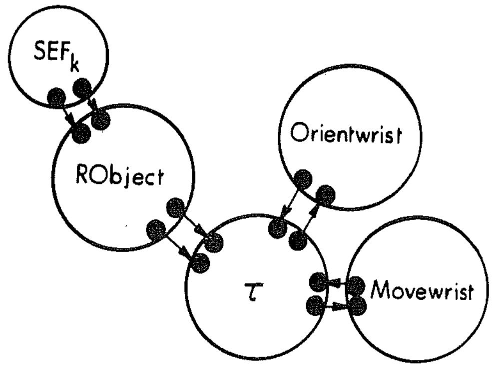

Figure A2.1. The Reach task-unit schema, with abbreviated syntax: Reach = [RObject − (MoveWrist, OrientWrist)].

tual schema SEF_k. The body of the SI then adjusts these values appropriately, and passes on the desired orientation and position of the wrist through output ports Ot and Pt. The coordinating schema τ then uses these values to direct the motor schemas Orientwrist and Movewrist to bring these values to the indicated positions. We now give the formal syntax. In what follows, the inserts enclosed thus ⟨ . . . ⟩ are not part of the task unit, but are rather interpolated comments to aid comprehension.

```
[Reach ($M_P$ $M_O$ $O_P$ $O_O$) (Pos Ornt)
```

⟨Sets up the ports for τ.⟩

```
    (RObject SEF)
```

⟨Specifies perceptual schemas.⟩

```
    (Movewrist Orientwrist)
```

⟨Specifies motor schemas. Next we instantiate RObject⟩

```
    (RObject(hlength, grot)
```

⟨with appropriate values of internal variables⟩

```
    ($O_P$←RObject(Pt), $O_O$←RObject(Ot),
    RObject(Pn)←$SEF_k(F_1)$, RObject(On)←$SEF_k(F_2)$)
```

⟨and port connections to τ and SEF_k. We then set up Movewrist, which has no internal variables, with its port connections to τ⟩

```
    Movewrist ()
    ($M_P$←Movewrist (actual), Movewrist (desired)←Pos)
```

⟨and similarly for Orientwrist.⟩

```
    Orientwrist ()
    ($M_P$←Orientwrist (actual), Orientwrist (desired)←Pos)
```

⟨Next we specify the local variables of Reach⟩

```
    (hlength grot Ptemp Otemp del)
```

⟨and then specify the body of τ, which cycles through the operations of reading the current position and orientation⟩

```
    (Ptemp:=$M_P$
    Otemp:=$M_O$
```

⟨and then issues commands to move and orient to the object (in steps specified by the local variable del)⟩

```
    IF Ptemp $\neq$ $O_P$ THEN Pos:= Ptemp+del ENDIF
    IF Otemp $\neq$ $O_O$ THEN Ornt:= Otemp+del ENDIF
```

⟨terminating only if orientation and position reach the desired values.⟩

```
    IF (Ptemp=$O_P$) AND (Otemp=$O_O$) THEN STOP ENDIF)]
```

Our second example illustrates the use of the FORALL statement. The schema Object? (Figure A2.2) initiates a parallel search for the presence of a given object in the environment. To simplify matters, we assume that the perceptual schema Objtest can decide the presence of the object from a set of four measurements reported by SEF (a different schema from that of the last example) over 4 output ports F_1, F_2, F_3, F_4.

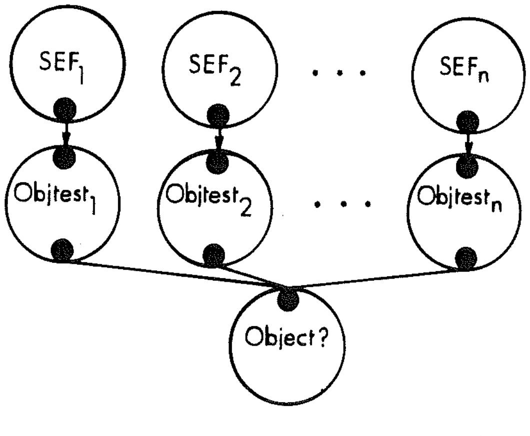

Figure A2.2 The parallel search initiated by Object?

```
    [Objtest ($F_1$ $F_2$ $F_3$ $F_4$) (judgment)
    ($DF_1$ $DF_2$ $DF_3$ $DF_4$)
```

⟨The input ports F_1, F_2, F_3, F_4 will receive the messages from the corresponding output ports F_1, F_2, F_3, F_4 of SEF, and DF_1, DF_2, DF_3, DF_4 are local variables that hold the desired values that characterize the object (we do not attempt to emulate here the subtlety of perceptual schemas of Appendix A!). The body of the schema thus makes the simple comparison⟩

```
    (IF ($F_1$=$DF_1$) AND ($F_2$=$DF_2$) AND ($F_3$=$DF_3$) AND ($F_4$=$DF_4$)
    THEN judgment:=1
    ELSE judgment:=0
    ENDIF
    STOP)]
```

⟨Object? then operates by instantiating copies of OBjtest for every region of the environment. We assume that an SEF SI is already in place to monitor each region. We also structure Object? so that it can feed information about the features of the located object to some task unit T.⟩

```
    [Object? (result) ()
    (T . . . . .
```

⟨This is shorthand for the specification of task unit T.⟩

```
    ($DF_1$ $DF_2$ $DF_3$ $DF_4$)
```

⟨Internal variables make the desired feature values available to each Objtest SI. The next instructions create one instance for each extant SEF SI:⟩

```
    (FORALL $SEF_i$ DO
    $Objtest_i$ ($DF_1$ $DF_2$ $DF_3$ $DF_4$)
    (Objtest($F_1$)←SEF($F_1$), Objtest($F_2$)←SEF($F_2$),
    Objtest($F_3$)←SEF($F_3$), Objtest($F_4$)←SEF($F_4$),
    Object (result)←$Objtest_i$ (judgment))
    ENDFOR
```

⟨These connections include a fan-in so that Object? will be bombarded by the judgments of all the Objtest SI for all the regions. Object? waits for a match, then sets up the task unit and terminates.⟩

```
    IF result=1 THEN
    T . . . . . .
    STOP
    ENDIF)]
```

The FORALL construct here expresses the inherent parallelism of a search. Of course in any actual implementation, be it in brain or robot, a key question will be to emulate this parallelism on a limited set of processors. For example, peripheral vision in humans can extract only crude feature properties; eye movements then direct foveal processing on the basis of these peripheral cues as well as on the basis of the current schema assemblage representing the environment (Didday & Arbib 1975).

The observant reader will note that this appendix describes a formal programming language that does not have *explicit* constructs for two concepts central to the vision schemas of Appendix A – namely, for the distinction between goal, hypothesis, and schema instance, and for the maintenance of activation/confidence levels. These can in fact be handled by introducing goal schemas and hypothesis schemas (the "request" and "contract" links will be built by the instantiation processes of Figure A1.2), and by maintaining the level as an internal variable to be updated by the body of the schema. However, it is a question for our current research to determine the best "marriage" of these two approaches to building a schema environment for efficient implementation on a network of computers.

APPENDIX C:

Matrix model of tectal–pretectal interactions. Appendix C is adapted from Appendix A of Cervantes-Perez (1985) and makes explicit the model of tectal–pretectal interactions used in Section 4. The input to the model comprises three arrays, R2, R3, and R4, which are generated by a model of the retina and represent the firing rates of three populations of retinal ganglion cells. The tectum–pretectum model comprises six layers, each of which is characterized by two arrays: the first, labeled with a lower-case letter, represents membrane potentials; the second, bearing the corresponding upper-case label, represents firing rates. There are five layers in the tectum – glomeruli (gl, GL), large pear-shaped cells (lp, LP), small pear-shaped cells (sp, SP), stellate neurons (sn, SN), and pyramidal cells (py, PY) – but there is just one layer (tp, TP) of cells in the current model of the pretectum.

The cells interact only via their "firing rates," which depend solely on each cell's own membrane potential. Thus in this model we do not need separate arrays to hold "old" and "new" values, but instead cycle through two steps:

Step 1: Using current values of firing rate arrays, compute new values of membrane potential arrays.

Step 2: Using new values of membrane potential arrays, use componentwise operations to form the new values of the corresponding firing rate arrays.

Step 1. Updating the membrane potentials

The difference equation: The membrane potential of each cell is described by a differential equation of the form

$$\tau_m \, dm(t)/dt = - m(t) + S_m(t)$$

which we will replace by the difference equation

$$\tau_m[(m(t + \Delta t) - m(t))/\Delta t] = - m(t) + S_m(t)$$

which yields

$$(1)\ m(t + \Delta t) = (1 - \Delta t/\tau_m)m(t) + (\Delta t/m)S_m(t).$$

The connection matrices: The influence of one layer, a, on another, b, will either be pointwise, in which case the strength of the connection will be written $w_{a.b}$, or some fan-in will be involved, in which case we represent the connections by a matrix $W_{a.b}$. In the present model, each matrix is 3×3, and is constant save for its central element. In what follows we shall use the abbreviation W{x,y} for the matrix

$$\begin{bmatrix} y & y & y \\ y & x & y \\ y & y & y \end{bmatrix}.$$

It should also be noted that the use of pointwise connections is unlikely to survive when the model is expanded, so that in what follows each term of the form $w_{a.b}.a$ has been replaced by $W_{a.b}*a$ where $W_{a.b} = W\{w_{a.b},0\}$.

With this we present the details for each cell type. In each case, we use the difference equation (1), so it only remains to specify τ_m and S_m for each choice of m:

Glomerulus

Here

$$\tau_{gl} = 2.3$$

and

$$S_{gl} = W_{r2.gl}*R2 + W_{lp.gl}*LP + W_{sp.gl}*SP$$

where

$$W_{r2.gl} = \{6.7,\ 0\},\ W_{lp.gl} = W\{8,\ 5.3\} \text{ and } W_{sp.gl} = W\{0.7,\ 0.7\}.$$

Large pear-shaped cell

Here

$$\tau_{lp} = 0.3$$

and

$$S_{lp} = W_{gl.lp}*gl + W_{r2.lp}*R2 + W_{sp.lp}*SP - W_{tp.lp}*TP - W_{sn.lp}*SN$$

where

$$W_{gl.lp} = W_{r2.lp} = W\{1,\ 0\}$$
$$W_{sp.lp} = W\{0.8,0.6\},\ W_{tp.lp} = W\{0.1,\ 0\}, \text{ and } W_{sn.lp} = W\{8.0,8.2\}.$$

Small pear-shaped cell

Here

$$\tau_{sp} = 0.9$$

and

$$S_{sp} = W_{r2.sp}*R2 + W_{gl.sp}*GL - W_{tp.sp}*TP - W_{sn.sp}*SP$$

where

$$W_{r2.sp} = W\{1,\ 0\},\ W_{gl.sp} = W\{1.0,0.5\},$$
$$W_{tp.sp} = W\{0.1,\ 0\}, \text{ and } W_{sn.sp} = W\{20.0,\ 0\}.$$

Stellate neuron

Here

$$\tau_{sn} = 1.6$$

and

$$S_{sn} = W_{lp.sn}*LP$$

where

$$W_{lp.sn} = W\{5.2,5.2\}.$$

Pyramidal neuron

Here

$$\tau_{py} = 0.12$$

and

$$S_{py} = W_{r2.py}*R2 + W_{r3.py}*R3 + W_{r4.py}*R4 + W_{sp.py}*SP + W_{lp.py}*LP - W_{tp.py}*TP$$

where

$$W_{r2.py} = W\{3.5,\ 0\},\ W_{r3.py} = W\{0.3,\ 0\},\ W_{r4.py} = W\{7.0,\ 0\},$$
$$W_{sp.py} = W\{2.0,0\}$$
$$W_{lp.py} = W\{0.7,\ 0.56\}, \text{ and } W_{tp.py} = W\{0.9,\ 0\}.$$

Pretectal cell

Here

$$\tau_{tp} = 0.02$$

and

$$S_{tp} = W_{r3.tp}*R3 + W_{r4.tp}*R4$$

where

$$W_{r3.tp} = W\{0.3,\ 0\} \text{ and } W_{r4.tp} = W\{5.0,\ 0\}.$$

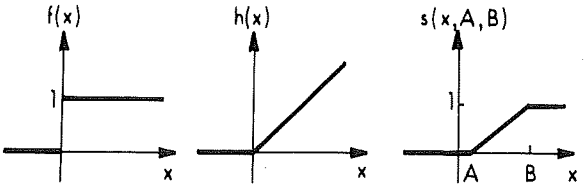

Figure A3.1. Thresholding functions.

Step 2. Updating the firing rates

The "firing rate" of a cell is obtained by passing the membrane potential through a "thresholding function" of one of the types shown in Figure A3.1:

$f(x) = \text{if } x \geq 0 \text{ then } 1 \text{ else } 0$

$h(x) = \text{if } x \geq 0 \text{ then } x \text{ else } 0$

$s(x,A,B) = \text{if } x \leq A \text{ then } 0 \text{ else (if } x \geq B \text{ then } 1 \text{ else } (x - A)/(B - A))$.

The updating of the firing rate is then accomplished simply by applying a thresholding function componentwise to the corresponding array of membrane potentials.

Glomerulus Here we simply take

$GL = gl$

but we still need two distinct arrays because of our updating convention.

Large pear-shaped cell	$LP = f(lp - 1.0)$
Small pear-shaped cell	$SP = f(sp - 2.0)$
Stellate neuron	$SN = h(sn - 0.2)$
Pyramidal neuron	$PY = s(py, 2.3, 5.0)$
Pretectal cell	$TP = h(tp - 3.8)$.

ACKNOWLEDGMENTS

Many of the models reported herein were developed while the author was at the University of Massachusetts at Amherst, in collaboration with R. Lara, F. Cervantes-Perez, A. S. Cromarty, D. House, T. Iberall, and D. Lyons, with partial support from NIH grant NS14971. Preliminary versions of this paper were presented at the symposia "Construction and Testing of Models of CNS Action" at the 29th Congress of the International Union of Physiological Sciences, Sydney, Australia, 1 September 1983; "Theoretical Trends in Neuroscience," sponsored by the Society for Mathematical Biology, at the 68th Annual Meeting of the Federation of American Societies for Experimental Biology, Saint Louis, Missouri, 3 April 1984, and "From Neuron to Thought," held in Florence, Italy, 29 March 1985.

NOTE

1. A series of papers by Pellionisz and Llinás (e.g., 1982) introduced a tensor theory of cerebellar function. Arbib and Amari (1985) have shown that the central claim of the theory – that the input to the cerebellum is a covariant intention vector transformed by a metric tensor encoded in the cerebellum to a contravariant execution vector – has not and probably cannot be substantiated. Later papers have made it clear that the work of Pellionisz and Llinás uses no tensor analysis beyond linear algebra. However, the work has had considerable *heuristic* value for a number of investigators (e.g., Gielen & van Zuylen 1986), encouraging them to seek an explicit analysis of coordinate transforms in the brain. It must be stressed that the analysis is unrelated to metric tensors, seeking rather to explore the utility of linear algebra (e.g., the Moore-Penrose generalized inverse; Pellionisz 1984) for the study of coordinate transforms, a method already applied in robotics.

Open Peer Commentary

Commentaries submitted by the qualified professional readership of this journal will be considered for publication in a later issue as Continuing Commentary on this article. Integrative overviews and syntheses are especially encouraged.

Biologically applied neural networks may foster the coevolution of neurobiology and cognitive psychology

Bill Baird

Department of Biophysics, University of California at Berkeley, Berkeley, Calif. 94720

Arbib's work and expertise span in detail an exceptional range, from experimental neurobiology through theoretical neural modeling to artificial intelligence and cognitive psychology. The program of synthesis across levels he describes and illustrates in his target article is worthy of special attention. The strategic importance in this program of the detailed neural modeling of a specific well-chosen physiological system as opposed to general analysis of abstract neural networks is strongly delineated in Arbib's discussion. There is a clear message here for the burgeoning field of "connectionism" and neural networks for associative memory. The strongest impact and development of these theoretical notions may be achieved where they are confronted with the real physiology of a particular system and refined to provide the bridge between biology and schemas of behavior that allows the cross-fertilization of structural and functional descriptions to occur.

Theoretical work in Freeman's laboratory at Berkeley on olfactory pattern recognition in the rabbit olfactory bulb (Freeman 1975) has proceeded by an interactive process like the incremental modeling approach Arbib describes, and we basically endorse his perspective. [See Skarda & Freeman: "Brains Make Chaos in Order to Make Sense of the World," *BBS* 10(2) 1987.] In this commentary I will focus on the issue of associative network implementations of perceptual schemas (an issue that is close to our work), expanding on selected points. Through an 8×8 array of electrodes monitoring spatial patterns in the EEG, we have observed the intermediate level of large-scale neural network activity during behavioral conditioning; this has given rise to a rewarding dialogue between experiment and theories of associative memory (Baird 1986a). The demands of this system have forced us to create and analyze pattern recognition networks that bifurcate from a single chaotic background state into a state space with multiple periodic and chaotic attractors (Baird 1986b). We hope that nature is instructing us to go beyond the systems that engineers are at present comfortable with so that we may discover new principles in these unusual features given to us by the natural design.

Although the "behavior" we study is a simple licking response to water reward, and we have not yet been motivated to use schema level descriptions, we are aware of the importance of the "action-sensation" cycle in perception and have an interest in visualizing the full sensory/motivational/motor control process as a system of neural networks (perhaps in "layers" of interconnected cortical areas). It was something of a disappointment to find that, despite Arbib's extensive work in motor control elsewhere, and the well-developed schema-level descriptions of approach behavior in frogs in his target article, there was no discussion of a corresponding neural model of motor or motivational structures or of the systems mapping cooperatively discriminated sensory features to motor output to implement the schemas described. There is mostly topographic mapping and local neighborhood interactions in the model presented, characteristic of a feature extraction or preprocessing operation, which in our system is followed by an associative memory type of network with global cross-correlating interconnection for object recognition. Might there be the need for such a system in the frog? It seems that relatively global (nontopographic) cross-mapping might be required at least to map ratios of features or patterns arising from identified objects to behaviorally related but otherwise structurally unrelated ratios of motor activation.

An important role of the intermediate level network model noted by Arbib is to guide and interpret experimental investigation where complex cooperative behavior of neurons is involved. This is certainly the case for associative networks. Single-cell "receptive field" or "trigger feature" investigations can be highly misleading (or uselessly incoherant, as they are in our system) when the network is cross-correlating its inputs. Apparent receptive fields may exist in such a net when the experimental situation is such that it always relaxes to the same attractor. Appropriate variation of the experimental context, however, could reveal a large variation in the apparent receptive field (as reported in some studies) when a different attractor is evoked and the relation of unit activity to input is completely altered.

Arbib points out that Marr's "computational" level or the functionalist's algorithmic level do not have the generality and a priori independence from biological implementation that a cognitivist might imagine, since different organisms may use different algorithms in such a way that the concept of the task itself is different. Thus the discovery of the implementational details of a biological system may contribute to the way a task is conceived and described at the functional level. Furthermore, the structural network level of description may be viewed as providing a compact way of encompassing large classes of symbolic algorithmic descriptions of the behavior of a system in the case where a network with a particular learning algorithm is capable of generating many specific algorithms to solve large classes of problems. Without knowing what specific symbolic description is appropriate we can know that the adaptive system is capable of solving the problem. John Holland's classifier system (1975), which uses a "genetic" learning algorithm, is perhaps a good example of an autonomous self-programming system. It works with binary strings whose symbolic meaning might be hard to interpret, yet succeeds in generating, evaluating, organizing, and utilizing new productions to perform the tasks required to meet its internal needs in an arbitrary external environment.

The centrality of instantiations

John A. Barnden
Computer Research Laboratory and Computer Science Department, New Mexico State University, Las Cruces, N. Mex. 88003

I applaud Arbib's incorporation of multiple, coexistent instantiations of schemas as a central feature of his framework. Not enough attention has been given in neural net and connectionist research to multiple instantiations of schemas, frames, situations, scripts, and so forth. In fact, insofar as Arbib's target article is about the general nature of a schema framework, I see its main substance as lying in the notion of multiple instantiations. If Arbib's general view of schemas did not encompass multiple instantiations, then one would presumably view schemas themselves as relatively permanent, interacting, concurrent processes. There is nothing particularly novel about the individual abstract schema programs Arbib proposes, and the general idea that information processing in the brain is based on interacting processes is hardly new. This, compounded with Arbib's understandable desire for a general framework rather than a particular specification language for schemas (see sec. 2, para. 1), means that without multiple instantiations we would not be left with very much in the way of a substantive proposal about abstract schemas in general. On the other hand, the inclusion of multiple instantiations presents deep, revealing, and approachable problems when the task of implementing the schema level in neural circuitry is considered.

I do not mean that multiple instantiations exhaust the substance of Arbib's target article in all its aspects. First, I am not impugning the intrinsic interest of the neural-level modeling efforts he reports. Second, the notion of cooperative computation would be important and interesting even in the absence of multiple schema instantiations. There are problems to be tackled in determining the nature of communication among cooperating agents and the means by which agents absorb incoming information. These are pressing issues in computer science, artificial intelligence, and cognitive science (including connectionism). In the Section 5 depth-finder example, the two cooperating subsystems (one disparity based, one accommodation based) can, I take it, be viewed as unique instantiations of two different schemas. This uniqueness does not mitigate the worth of the model or the schema notion underlying it.

I also find Arbib's plea for an intermediate level between behavior and neurons a refreshing corrective to the lack of such a level in much neural-net and connectionist work. Quite apart from the point that intermediate levels are a heuristic aid in managing a large research effort, they can also free one to discover useful reductions of high-level notions to low-level ones that could otherwise escape one's attention – through being, perhaps, too complex or deviant to conceive of unitarily.[1]

Arbib is right to emphasize multiple schema instantiations, but I wish that he had pursued the matter further. The depth-finder example of Section 5 is prominent in the paper, yet it does not involve multiple instantiations. The other examples, in the main text and Appendix B, do not involve *detailed* attention to multiple instantiations, and there is no consideration of how multiple instantiations would be neurally realized. As I shall now discuss, the multiple-instantiation issue has relationships to longstanding issues in brain theory and connectionism, and poses interesting problems for these fields.

One old problem in connectionist research is that of avoiding "cross-talk" between different pieces of information. For example, if one supposes that a connectionistically implemented agent is simultaneously entertaining the ideas that John loves Mary and Bill loves Sally, one has to take a certain amount of care to ensure that the agent's internal state is not also one that would obtain if the propositions were instead that John loves Sally and Bill loves Mary.[2] A standard approach to the issue is based on recruiting neurons or neural assemblies to represent the particular instances of a situation class (e.g., loving) as well as having a neural assembly standing for the situation class itself (see Hinton 1981). Recruitment is also advantageous in accounting for the ability of an agent to entertain any novel, short-term, complex propositions (or data structures) in the first place. The notion of temporary, short-term recruitment in turn raises such questions as, for example: How is it managed? Does it rely on synaptic-weight change or on some other mechanism? How can it be made fast enough? How economically can inferential and other information-processing mechanisms respond to neural structures involving recruited – and therefore in a sense unpredictable – assemblies? How are recruited assemblies demobilized? – a nontrivial question at least in the case of "distributed" connectionist systems in which a given assembly can share many neurons with other assemblies. The report of a recent workshop on connectionism (McClelland et al. 1986) identifies cross-talk and the more general issue of accounting for possibly novel, complex temporary data structures as being of major concern. The relevance to the multiple-instantiation issue is clear – simultaneous presence of several propositions about loving is similar to, or even an example of, the simultaneous presence of several instantiations of an Arbib schema.

Suppose for definiteness that a schema S is neurally realized as some particular neural net N. What then does it mean to say that an instantiation of S is present? If only one instantiation at a time were ever allowed, an instantiation's existence could simply be a matter of certain state parameters of N having certain values – e.g., certain neurons having certain firing rates. But in the multiple, coexisting instantiation case, things are not so simple – e.g., we cannot suppose that a neuron can simultaneously be firing at several different rates. We can ask general questions such as: Can we produce a theory in which different instantiations are somehow superimposed states of N itself? Or are the different instantiations different copies of N in some sense (these copies being dynamically recruited, or perhaps permanently existing but only intermittently active)? Or do the different instantiations involve neural networks that are not copies of N but are temporary holders of different sets of parameter values for S and different "program counters" for S (so that the networks might be copies of only a part of N)? If instantiations involve neuronal recruitment, what are the answers to the above questions about recruitment? How are the answers to all these questions affected by a consideration of the number of instantiations of a given schema that can be simultaneously present?

This discussion of instantiations can be modified to account for the possibility that it is only instantiations, and not schemas themselves, that are identifiable with neural circuits. A schema could, for instance, be construed merely as a propensity to create or activate neural circuits of a certain form that act as instantiations.

A special problem concerning data communication among computing agents arises for a brain-oriented framework like Arbib's. There are several different types of neural encoding that might be envisaged for the data. For example, on one channel of communication the information might take a retinotopic form; on another, it might be encoded in the firing rates of some neurons with no special organization. I appreciate that at the schema level Arbib may not wish to take such factors into account because they are implementation details (although he is sensitive to them, as point 2(b) in section 6 shows). At the same time, I suspect that such factors should be allowed to affect abstract schema formalisms. For example, a particular type of encoding (perhaps by single-neuron firing rate) may have a large inherent imprecision, so that only data channels that could tolerate that imprecision should be so encoded, whereas other types of encoding (e.g., fine-grained retinotopic encoding of position information) may involve much less imprecision. It might therefore be advantageous to specify, in a schema formalism itself, the degrees of precision that various channels (ports) require. This would (a) provide a guard against uncritically allowing schemas to transfer values on low-precision channels to high-precision ones. More generally, it would (b) help to ensure that the behavioral part of a schema does not involve excessive hidden conversion between different sorts of encoding. It would also (c) usefully constrain the types of encoding allowed in implementing the schema system at hand in neural circuitry. Points (b) and (c) suggest that a schema formalism could even allow particular encoding techniques to be specified or suggested for channels.

There is another enhancement to the schema framework that I would like to propose and that I suspect would be congenial to Arbib. Appendices A and B imply that the behavioral part of a schema is to be specified by means of a fairly conventional analogue to current programming languages (although the imprecise description of schemas in the main text leaves the door open to other sorts of behavior specification). The suggested enhancement is to allow the behavior to be specified instead by means of a mathematical description of an input–output function linking input ports to output ports, or, more generally, of a relation (in the set-theoretic sense) linking the ports. The introduction of a *program* to compute the function or relation could then be part of the task of *implementing* the schema in lower-level terms; and it may not even be appropriate to think of using a program (as opposed to, say, a connectionist subsystem with no convenient abstract algorithmic characterization) on the implementation route. I would in fact favor a hybrid approach in which both programs and mathematical descriptions are available as behavior-specification tools.

NOTES

1. In my own research on connectionist models for complex short-term inferential information processing I have eschewed the typical idea of thinking of *abstract* connectivity among pieces of information as being mapped in any direct way onto *hardware* connectivity among units in a network. Instead, the mapping appeals to an intermediate level in which there is a notion of relative position of data items in certain representational media. This relative position at the intermediate level manifests itself in terms of hardware connectivity in a complex way (see Barnden 1985; 1986).

2. In some contexts, however, one might want to claim that cross-talk effects accurately reflect human information processing. This approach is taken in a limited visual-processing context by Hinton and Lang (1985).

Advantage of modeling in neuroscience

J.-P. Ewert

Department of Neuroethology, University of Kassel, D-3500 Kassel, Federal Republic of Germany

Professor Arbib's successful approach to *Rana computatrix* follows the cybernetic enterprise of addressing problems of sensorimotor coordination by constructing model systems. Because there are no boundaries between biological, artificial, and hybrid systems in cybernetics (the "bridge between sciences"), cross-disciplinary collaboration makes comparative studies in brain function feasible and testable. The amphibian brain can be regarded in a certain sense as a kind of "microcosm" of the highly complex primate brain, as far as certain homologous structures and assigned ballistic preplanned and pre-programmed target-oriented behavioral processes are concerned (e.g., see Bullock 1983; Northcutt 1986). It is thus possible to investigate in amphibians comparable visuomotor functions without having to face mammalian complexity. During our collaborative research with systems theorists we have learned that progress in the understanding of neural principles proceeds with guidance from model networks. Modeling – paralleled by simulations – transforms theory derived from data into testable structures. Where empirical data are difficult to obtain because of structural constraints, the modeler can provide assumptions and approximations that, by themselves, are a source of hypotheses and subject to predictions. If a neural model is tied to empirical data, it can be used to predict results and to stimulate new experimental tests whose data may in turn improve the model. The algorithms derived from the model have heuristic value for comparative studies in neuroscience and can also be of applied scientific interest. We will select some examples to illustrate the advantage of modeling.

Predictions. Our first systems-theoretic model explaining the neuroethological data of the toad's prey selection in response to moving visual objects (Ewert & von Seelen 1974) was based on interacting homogeneous networks coupled with lateral excitation and lateral inhibition and was characterized by the following properties: (1) Specific filter operations in space and time domains are carried out by retinotectal and retinopretectal networks; (2) stimulus discrimination is a result of subtractive network interactions; (3) the output of the model network is mediated by prey-selective tectal neurons that decode the "prey schema" (encoded in terms of circuits provided by functional units); (4) pattern recognition can be implemented with this circuitry by comparing (cross-correlating) the visual input with its filter property (see also Suga 1984). Due to symmetrical network couplings, such a model predicts that the preference of a stripe moving in worm (W) configuration versus the same stripe moving in antiworm (A) configuration should be *invariant* under the direction of object movement in the x-y coordinates of visual space. Experimental tests of the toad's prey-catching (Ewert, Arend, Becker & Borchers 1979) and the response of prey-selective tectal T5(2) neurons (Ewert, Borchers & von Wietersheim 1979) have shown movement directional invariance of the W/A configural preference. The invariant ("universal") refers here to invariant relationships provided by the neuronal circuitry.

Limitations. Without simulations it is often hard to realize that a given model network actually includes many different models, depending on the choice of network parameters. This can be illustrated by the following example. Since elongating an antiwormlike stripe perpendicular to the direction of movement only changes spatial parameters whereas elongating a wormlike stripe parallel to the direction of movement influences both spatial and temporal components of the stimulus, the W/A (worm/antiworm) discrimination (Ewert 1969) can be explained in part by an "asymmetry in the time domain" despite symmetrical network couplings: The worm darkens a retinal area longer than does a corresponding antiworm. Computer simulations by model networks (Cervantes-Perez 1985) therefore predict – for specific parameter settings – that changes in the dynamic domain should, by variation of the visual angular velocity of the stimuli, alter the stimulus selectivity of the network's output and, at an extreme, even lead to an inversion of W/A preference. However, the experimental data on the common toad's prey-selective T5(2) neurons (see Ewert et al. 1978, p. 44; cf. also the corresponding "H>S>V neurons" by Roth and Jordan 1982, p.395) and prey-catching (Burghagen 1979) have shown that the W/A configural preference is *invariant* under changes of stimulus angular velocity. (The velocity-dependent "inversion" of the fire salamander's W/A preference reported by Luthard and Roth [1979] has unfortunately not turned out to be reproducible – Roth, personal communication; Himstedt 1982; Finkenstädt 1983 – and can only be explained by experimental error.) Although state-dependent processes modulate certain characteristics of visual neurons so as to influence object discrimination (related to stimulus area and shape) and to shift category boundaries, an inversion of the W/A configural preference in prey catching could not be induced, even by specific training experiments (Ewert 1981; for definitions of "configuration" and "shape" see Appendix in the target article by Ewert, this issue). The reason for velocity invariance in the W/A preference is obviously linked to (1) the different properties of tectal W-coding and pretectal A-coding circuits and (2) their subtractive interactions. This example shows that the potential of a model system may be much greater than the actual requirements of the brain. The interaction between theory and experiment thus allows one to determine the operational range of the validity of certain parameters of the model.

Approximations. The models of Arbib and his coworkers proceed at two levels: schemas and neural networks. Arbib's concept of schemas extends the usage in classical ethology. Whereas the classic work of Jacob von Uexküll (1909), Niko Tinbergen (1951), and Konrad Lorenz (1943) studied behavior in terms of the release of single motor schemas (fixed action patterns) by the activity of separate perceptual schemas (innate releasing mechanisms), schema theory allows us to study the interpretation of the environment in terms of the *interactions* of perceptual schemas, yielding appropriate coordinated control programs for motor schemas and compound motor coordinations (cf. "interacting functional units" of Figure 19C and "sensorimotor codes" of Figure 22 in Ewert's target article). The model of detour behavior is a good example of this. But since Arbib and his coworkers seek neuronal implementations of schemas, they also share in the advantage that neuroethology has over classical ethology: We can come to understand how the evolution of the brain constrains the algorithmic specification of schemas (see also Ewert's commentary on Anderson's target article, this issue). The method of incremental modeling ("evolving family of models") similarly gives insight into brain mechanisms by studying how to change the model to accommodate the function of more and more schemas and their interaction.

These models incorporate properties concerned with space- and context-dependent variables, as well as coordinated control programs that link perceptual and motor schemas. Investigating sensorimotor codes we are faced with the question of how these might "unlock" specific bulbar/spinal motor pattern generators (MPGs), the motor schemas (see Figure 22 in Ewert's target article). Our recent discovery of bulbar neurons displaying multiple properties (e.g., involving T5(4), T4, and cyclic bursting characteristics) suggests that interactions among highly integrative cells are incorporated in a relatively compact "reticular" network (Ewert et al., submitted; Schwippert & Ewert, submitted). Since in the interwoven structures of the reticular system functions are not discretely localized but integrated (which is a problem for lesion studies), guidance from model networks may be particularly fruitful here in order to investigate coordinated control programs and understand MPGs (Selverston 1980). Concerning the "interwoven" nature of the bulbar reticular structure one can ask whether MPGs for various rhythmic action patterns depend on a general bulbar pattern generating circuitry that (as a function of specific inputs from different "command releasing systems" or sensorimotor codes) develops appropriate patterns of excitation and inhibition. Alternative models may involve individual, partly interacting, oscillators.

Applications. The study of biological sensorimotor systems holds the promise of providing elegant solutions to practical problems in technical devices. Natural pattern recognition systems, for example, successfully deal with different invariance conditions (see commentary by Stevens on Ewert's target article, this issue). Comparable properties have proved difficult to implement in the optoelectronics of machines (e.g., position/orientation/velocity invariance in the recognition of work pieces on assembly lines; prevention of occlusion by background structures). Computer simulations of natural information-processing systems represent operating algorithms, derived from the basic properties of the model, which could be implemented in robotic systems.

We hope that the current dialogue between *Bufo bufo* and *Bufo computatrix* will help solve some of the problems in our understanding of adaptive sensorimotor and motor pattern generating processes.

Levels of psychological reality, Arbib's "schemas," and matters maybe metaphysical

Keith Gunderson

Department of Philosophy and the Minnesota Center for Philosophy of Science, University of Minnesota, Minneapolis, Minn. 55455

I would like to raise a few noncontentious methodological and perhaps quasi-metaphysical questions about matters only inconspicuously untouched in Arbib's target article. Arbib hopes his model will help reveal principles of neural organization and will mediate between that level and complex behaviors that, I assume, are in some way psychological. I take it that "schemas" are a kind of conceptual dowry donated to a hopeful marriage between functional cognitive levels of description and those involving neural localization and implementation in neural circuitry.

I am not competent to assess the details of the latter levels, but I am impressed by arguments of Cherniak (1987) that various idealizations of neural (cognitive) function do not seem to accord with neuroanatomy (one requiring a brain as "big as a bathtub") and so on. So Arbib's modeling of schemas "by interacting layers of neuronlike elements, or by nets of 'intermediate-level' units" which are "little constrained by anatomy or physiology" may be more risky than it appears. But setting that aside, how is the top (behavioral) level to be circumscribed?

Arbib claims that his schemas (simultaneous computing agents) "must be 'large' enough to allow an analysis of behavior at or near the psychological level." What, however, is the difference between a psychological level and a "near"-psychological level insofar as the explanatory target of *Rana computatrix* is concerned? Unless we know this, exactly what Arbib's "large"-enough schemas initially apply to at the top tier of analysis remains unclear. Let me label this the Psychological Level Triangulation problem (or the PLT problem).

A PLT problem arises for any psychological research program that purports to model, explain, reduce, or eliminate methodologically or metaphysically puzzling mental phenomena at some "deeper" – clearly physicalistic (organismic) – level of analysis. I also think the PLT problem pops up in the strategic framework espoused by Arbib, which does not attempt to reduce cognitive science to one fundamental level but instead uses what he calls a " 'two-way reduction' whereby the reality of different levels may be recognized and reconciled to create a richer whole." Related to the foregoing is the question of what representational role and psychological reality Arbib's schemas assume in his overall conceptual picture.

Schemas are said to be "approximations to reality" – both a psychological and a neurophysiological reality, I would assume. But what kind of *psychological* reality? I would suppose that the schema models of *Rana computatrix* are designed to provide a base from which to generalize (at least with respect to method) to how one might dissect human visual-cognitive competences. Yet one must decide whether such competences are to be characterized solely in terms of behaviors and their underlying neurophysiological causes or also in terms of whatever it is (introspectively) like to make and be aware of making various perceptual discriminations. Any interest in such a decision arises only if one focuses on the extrapolation from the frog–toad model to a human one, since not being a frog or a toad, one is not faced with the possibility of using introspectively accessed data. But in the case of human beings there is always the suspicion, no matter what the methodology, that just such data do in one way or another play some role, whether overtly or covertly. For example, in Newell and Simon's (1963) early computer simulations the use of verbal protocols in constructing the program was an overt acknowledgment of the possible relevance of introspective data. It constituted one kind of coping with what I have called the PLT problem. Typically, however, I think the role of such first-person psychological data, if any, is much less obvious. In connection with Arbib's project this topic concerns the psychological reality of his schemas.

Given that schemas are used to represent cognitive-perceptual phenomena, how are they to be interpreted vis-à-vis the introspectively accessible strata of such phenomena? Do any of them represent *any of that?* or do none of them make "contact" at such levels? Similar considerations arise for Marr's (1982) approach to vision. In discussing the "range of perspectives that must be satisfied before one can be said, from a human and scientific point of view to have understood visual perception" Marr wrote: "First and foremost, there is the perspective of the plain man. He knows what it is like to see, and unless the bones of one's arguments and theories roughly correspond to what this person knows to be true at first hand, one will probably be wrong. . ." (pp. 4–5). But how is that first-person point of view of the "plain man" to be represented? Is it only a rhetorical starting point that gets lost in the model-building shuffle?

Arbib's claim that his "schemas become 'more real' as their functional analysis is refined into assemblages/programs of subschemas that allow either a more subtle analysis of behavior or an improved mapping of function to neural structure" suggests that such "contact" – with "qualia," "raw-feels," or whatever one least dislikes calling them – is either irrelevant or at most highly covert. If so, is this because they seem ontologically undesirable? Methodologically Intractable? Or both?

Related to these issues is the question of whether the various explanatory schemas are to be viewed as being equally psychological (or "mentalized"). A schema representing "generic knowledge about some domain of interaction (e.g., a chair and how to sit on it)" at least initially suggests some (complicated and considerable) degree of understanding and intentionality. But is this true for schemas in general? (Might Arbib, for example, wish to borrow a controversial tactic from Dennett (1978b) and construe them as hierarchically "homuncularized" – schemas at each level performing diminishing amounts of cognitive donkey work yet allegedly accounting for our various intentionalized tasks without spiraling into a regress, or, at the end of it all, assuming some strange metaphysical status? Again, with respect to the work of Marr and others in computational vision, one can raise a similar question. Are the various levels of computation that allegedly underlie visual cognition to be viewed as equally or diversely psychological or, as could be the case, not psychological at all? Clearly they are not to be *dis*qualified as candidates for psychological reality simply because we lack introspective access to them. Yet one would like at some point to have a sense of how to distinguish between internal states and processes that are mental though not introspectively available, and those that are purely noncognitive physicalistic processes underlying them. An underlying cause of something psychological need not itself be psychological. And, of course, there are all sorts of information-processing subsystems in us capable of instantiating a myriad of computer programs, which are decidedly noncognitive in nature. (As Searle [1980] pointed out in *BBS*, there is a level at which the stomach can be described as an information-processing subsystem, but the stomach is not thereby endowed with any capacity for understanding.)

Perhaps these questions are not "quasi"-metaphysical, as I suggested at the outset, but floridly metaphysical, and simply a roundabout way of coaxing to the surface some suggestions as to how the mind–body relationship sorts itself out in Arbib's conceptual scheme of schemas.

The biotope of *Rana computatrix*

P. I. M. Johannesma
Department of Medical Physics and Biophysics, University of Nijmegen, 6525 EZ Nijmegen, The Netherlands

Rana computatrix, created by Michael Arbib and developed in his computational school, is intended "to show how one constructs an evolving set of model families to mediate flexible cooperation between theory and experiment."

In a Grimm's fairytale a prince is incarcerated in the body of a frog; in the story of Arbib a frog is encaged in the frame of a machine. Whereas Grimm introduces a beautiful princess to release the prince from his frog-type confinement, Arbib does not suggest such a happy ending for his frog: The machine appears to form a permanent residence.

In order for *Rana computatrix* to fulfill its role in the cooperation of theory and experiment it is desirable that frogs, natural and artificial, share a common biotope for perception and behavior. We share the view of Arbib that the conformity of *Rana computatrix* and *Rana naturalis* is not only determined by their structural similarity but at least as much by their functional congruence (see Arbib's Figure 1). Functional aspects are not clearly defined in Arbib's descriptions, however. As a consequence, *Rana computatrix* has no proper access to the biotope of *Rana naturalis.* The liberation of the computational frog from algorithmic schemas is the objective of this commentary. We will not be able to attain this goal completely, but we will try to indicate a way out.

For an arbitrary system, three aspects or levels of description are relevant: structure, process, and function.

Structure is the material substrate of the system; for many types of systems this may be considered as a set of elements and relations. In the nervous system these are the neurons and their axonal-dendritic connections.

Process is the dynamic sequence of states occurring in the structure. Parameters change on a long-time scale; variables change on a short-time scale. For an autonomous system, state variables give a complete description of the process. In the nervous system we can consider synaptic strength and threshold as parameters and the dendritic-somatic currents and/or potentials as state variables, with the action potentials as derived variables (Johannesma et al. 1986).

Function is the relation of the system with its environment. For most systems this relation is the transformation of input into output, stimulus to response. This class of systems we call transformational. If we know the transformation or mapping of an arbitrary input x into the associated output y, then we know the function of the system. However, there is a wider class of systems whose function is not only transformation but which includes representation. A representational system forms an internal model of its environment. Brains and computers belong to the class of representational systems. The visuomotor system, as part of the nervous system, transforms the image of a moving object into the rotation of the eyes following this object or the movement of the body with respect to this object, but it also constructs an internal model of the visual scene (even into the future) and of the possible behavioral attitudes with respect to this object. If a tree frog jumps at a fly in flight it is not responding directly to a physical stimulus but generating a behavioral program with respect to a perceptual construct derived from the physical stimulus based on internal state and a possible set of behavioral acts.

The distinction between process and function and between transformation and representation may clarify the concept and use of schemas in perception. Process is the transformation of the input of a schema into its output. For example, in the control program for reaching toward and grasping an object shown in Figure 4 of Arbib's target article the "VISUAL LOCATION" schema transforms "visual input" and "recognition criteria" during "activation of visual search" into "target location." On the other hand, the function of "VISUAL LOCATION" is to represent "target localization." Although "target localization" is coded in internal variables (e.g., a spatiotemporal pattern of action potentials), its meaning is the location of a target in external variables. The neural activity pattern represents the location of a fly in the physical world around the frog. In order to understand the meaning of the "neural represention of sensory stimuli" it may be useful to construct the "sensory interpretation of neural activity." If the neural representation is considered as a mapping P from sensory space into neural space, then the sensory interpretation is the inverse map Π of neural space into sensory space. This is illustrated in Figure 1.

Transformations P and Q can be divisioned in subtransformations or schemas (see, e.g., Figures 4 and 11 of Arbib), leading to an elaboration of Figure 1 and to a number of forward maps P_i and inverse maps Π_i. The dynamic development of the neural variables determined by brain structures and processes corresponds to the trajectory of movement in sensory space.

Our assertion is that the formal description of the nervous system should be in terms of structure, process, and transformation, that is, the forward mappings P and Q, while understanding the function of the nervous system should be based on the representation given by backward map Π and forward map Q.

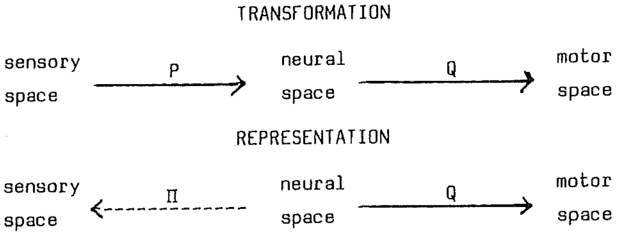

Figure 1. (Johannesma). The process of transformation and the function of representation of neural variables. The transformation diagram shows the flow of neural signals as a physiological process; the representation diagram indicates the function of neural activity as an internal model.

In sensory parts of the nervous system the backward map Π can be derived from the response of single neurons to sensory stimuli. The formulation of the experimental data is normally in terms of receptive fields of the neuron with spatio-spectro-temporal dimensions (Aertsen & Johannesma 1981; Eggermont et al. 1983; Johannesma & Eggermont 1983). The inverse map Π can be computed by the appropriate combination of forward and backward transition probabilities. An approximation is given by superposition of the receptive fields of the neurons that are active. Since receptive fields are defined in sensory space, the result of this procedure is a signal or a distribution of signals in sensory space. (For a mathematical description and some results of simulation see Johannesma 1981; Johannesma & Aertsen 1982; Johannesma et al. 1986.)

If sensory stimulus x induces the neural activity pattern z, be it deterministically or stochastically, then

$$z = P\,x.$$

If ξ is the sensory interpretation of neural activity pattern z, then

$$\xi = \Pi\,z.$$

As a consequence, the sensory interpretation ξ of a sensory stimulus x is

$$\xi = \Pi\,P\,x$$

where Π P is a deterministic or stochastic mapping of sensory space onto itself.

A demonstration of these ideas can be given in relation to Figure 23 of Arbib's target article. The quantity displayed in the graphs is the activity of the accommodation field and the disparity field. This activity forms an internal variable of *Rana computatrix.* In my opinion no appropriate metric has been or can be defined to compare internal variables of model and animal. Even neural variables of different animals cannot be compared directly; they form different languages with individual syntaxes; only in the semantic domain might a common frame of reference exist. Application of a sensory interpretation of the internal variables by an inverse map of neural space into sensory space would result in a structure of the visual space of *Rana computatrix* that could then be compared to the visual space of *Rana naturalis.*

Limits of perception, both in the natural and in the artificial frog, are given by the map π P. The evaluation of the quality of this map should be made by assessing the perceptual resolution of the animal by observing its behavior. We do not imply that the map π is available to the subject, natural or artificial, but in our opinion it allows an external observer to compare perceptual abilities of *Rana naturalis* and *Rana computatrix.* Along these lines artificial and natural frogs are not compared with respect to their internal variables, but as companions sharing a common biotope.

Structure and process in schema-based architectures

Pat Langley
Department of Information and Computer Science, University of California, Irvine, Calif. 92717

I was pleased to read Michael Arbib's argument for different levels of description in cognitive science; in the flurry of recent activity over connectionist theories of behavior, the higher levels of abstraction have often been overlooked, despite continuing advances at those levels. Newell, Simon, and Shaw's (1958) original arguments for the computer simulation of human behavior focused on the notion of *functional* descriptions, and Arbib reintroduces that concept in justification of his schema-based theory.

Arbib also argues for linking neural-level and functional-level explanations. Although I respect this goal, I am not convinced that this is the best path for cognitive science at this stage in its development. The history of physical science shows that higher-level theories can be successfully pursued independently of lower-level ones. Chemistry made steady progress for well over a century before the quantum theory and particle physics provided deeper explanations for its basic phenomena. Similarly, cognitive science has made steady progress at the functional level, developing symbol-based theories of problem solving, natural language processing, vision, memory, and learning that have provided us with a broad understanding of human cognition and perception.

Certainly, at some stage cognitive scientists will want to connect their functional models with neural models, and Arbib's work presents an excellent example of this approach. He argues convincingly that this can already be accomplished for simple instinctive behavior, such as that occurring in frogs. But applying this approach to the more varied and complex mechanisms that underlie human intelligence is another matter entirely, and it will be some time before we are capable of such feats. In the meantime, we should not ignore the mechanisms involved in problem solving and natural language processing simply because we cannot yet explain them in neural terms. Rather, we should continue to study them at the functional level, extending the significant strides that have already been made in this direction.

Now let me turn to the details of Arbib's schema theory. I firmly believe that our computational theories should be cast within an overall model of the cognitive architecture, so I was encouraged by Arbib's steps in this direction. Moreover, most research on schema-based architectures (e.g., Schank 1982) has focused on issues of representation and memory organization, dealing with process concerns in an ad hoc manner.[1] Given his interest in motor behavior, I hoped Arbib would break away from this trend, but when all was said and done, the only processes he had proposed were vague and rather standard notions of schema "activation." I believe that schema-based theories have much to offer, but I do not think we should attempt to map them onto neural models until they have been more completely specified at the functional level.

I was also disappointed that Arbib's schema framework did not make more contact with the well-established phenomena of cognitive psychology. He incorporates the notions of short-term and long-term memories, but he does not attempt to explain the well-known limits on the former or the occasional inaccessibility of the latter. He invokes the distinction between declarative and procedural knowledge, but then fails to elaborate on their connection. Are the retrieval mechanisms different for the two forms of knowledge? Are separate learning mechanisms involved? Well-defined structures are of little use for making predictions if the associated processes are ill defined.

I also felt the focus on parallelism among schemas was misplaced. Parallelism at the neural level does not imply that it exists at the functional level, despite its current popularity. There is certainly evidence for some parallelism at the functional level: It seems likely that much of retrieval and classification operate in this mode. However, there is also strong evidence for an attentional bottleneck that requires serial processing on tasks such as problem solving and parsing. Recent production-system models (e.g., Rosenbloom & Newell 1987) have addressed this dual character of the human information-processing system and I would welcome a concerted effort to handle these phenomena within a schema-based architecture.

Finally, I found the representation for motor schemas to be somewhat implausible. Artificial intelligence and cognitive science have succeeded largely because they rejected traditional computer science representations and developed their own formalisms for describing symbolic structures and processes. Arbib's schemas suggest an influence from the field of robotics, in which most researchers continue to use numeric representations and algorithmic control structures. I do not mean to suggest that motor behavior is best described using semantic networks or production systems; I fully expect that this domain deserves its own representational language, and that this will incorporate levels of detail about the body and the environment that existing AI representations could not easily support. But neither can I believe that Arbib's representation will give the flexibility and adaptivity required to model human motor behavior, even if it suffices for frogs. Nor can I see how it would support models of motor skill improvement and acquisition.

I am probably being too harsh. Actually, I am quite pleased to see someone seriously addressing the nature of sensorimotor intelligence; cognitive science has nearly ignored this topic, to its own detriment. I also respect Arbib's goal of linking functional-level accounts to neural-level models and encourage researchers interested in the interface between schemas and behavior to join him in this quest. But I encourage other researchers to work toward a more complete account of the human cognitive architecture at a functional level. Chemistry did not disappear with the development of atomic physics, and functional-level explanations will continue to play a central role in cognitive science as long as the field exists. We have already made significant strides in understanding the processes and structures that underlie human cognition, but we need to expand upon these initial results and integrate them into a coherent theory of behavior. This quest is fully as important as determining the mapping between neurons and functions.

NOTE

[1] In contrast, research in production-system architectures (Anderson 1983; Klahr et al. 1987) has focused on process issues in preference to problems involving memory and its organization.

Cognitive modeling: Of *Gedanken* beasts and human beings

Dan Lloyd
Department of Philosophy, Trinity College, Hartford, Conn. 06106

It's an article of faith, reflected even in the title of this journal, that top-down cognitive science and bottom-up neuroscience will ultimately converge on unified models of complex behaving systems. Arbib has gone a long way toward realizing that synthesizing ambition, and the questions pressed here are primarily invitations to further incremental development. *Rana computatrix* and other computer implementable models in behavioral neuroscience spark concern about two related issues: completeness and accuracy.

A complex neural model should be *complete* in two senses: The behavioral capacities we analyze and model from the top should bottom out in neurally plausible implementations; and the relevant pathways modeled should extend without break from sensory input to motor output. Otherwise, we risk smuggling in a homunculus, or at least an anuranculus – a black box that makes the model go but has no apparent correlate in the living system – and we thereby end up modeling our own handwaving. Though under construction, *Rana computatrix* comes close to satisfying this standard of completeness, but not every theme sounded in the target article has yet found a home in a complete model. The target article completes vertical links with respect to facilitation effects and prey discrimination, and alludes to other successes. But the current depth and detour detection schemas Arbib discusses (sections 3.3, 5) are mainly top-down analyses, with some support from lesion studies. Though something like these schemas will ultimately play a role in anuran psychology, as Arbib notes (sections 1, 3), the particular schemas proposed may not reduce to well-defined localized brain functions. I take it that schemas that turn out to be bottomless are thereby undermined. Thus the scaffolding from top to bottom is, at present, somewhat top-heavy.

The connectionist models familiar from Rumelhart and McClelland (1986) meet the standards of completeness, but Arbib rightly questions them for omitting intermediate levels of neural organization. This raises a second inevitable consideration in large-scale neural modeling: *accuracy*. *Rana computatrix* incorporates a range of neuroscientific data that are either unavailable or ignored in most PDP (parallel distributed processing) models. Thus Arbib's model is presumably closer to reality than the relatively undifferentiated architectures proposed by most connectionists. Nonetheless, there remain questions about accuracy. At several points (1.0, 4.1, 4.2), Arbib notes that producing the right characteristics required tuning the parameters of his model tectal columns. Appendix C indicates the numerous parameters there to be tuned. One wonders what guides the choices of adjustments among these parameters. Furthermore, in simplifying the physiology of the tectal column (Figure 2), and in assembling the synthetic tectum (4.2), the modeler risks ignoring a physiological determinant of behavior with nonobvious but real effects. In dynamic systems like the nervous system, an incremental departure from accuracy in any of these details can have ramifying effects, requiring the equivalent of compensatory epicycles to keep the model on track. Increments have more than incremental consequences.

In the ideal world, models would be maximally complete and perfectly accurate. But in practice, completeness and accuracy force the researcher into trade-offs. On the one hand, constructing a complete model demands some educated guessing; on the other, at this point our accurate knowledge is limited to parts of all but the simplest systems. Which way should the modeler build? I suggest that, if one must choose, *completeness* is the most useful goal. Arbib points out two reasons for this: Focusing only on components can lead one to overlook emergent effects of cooperative computation; and working with complete models, even if wrong in detail, nonetheless provides us with analytical tools applicable to future data and future, more accurate, models (sections 3, 6). I think the reasons to push for completeness go beyond these, however, and reveal a further source of the value of Arbib's work for cognitive science.

In a commentary in the first volume of this journal, Daniel Dennett (1978a) described a research strategy for understanding human cognition that surely smolders (though perhaps not by Dennett's spark) in the minds of most students of animal cognition. He proposed that one approach the complexity of humans by looking at simpler systems, first solving cognitive problems as they arise in these "simple minds" and then bootstrapping toward increasingly complex and humanlike cognizers. Two sorts of systems appealed to Dennett as fruitful stepping stones: living systems and artifactual systems born of engineering imagination, "Martian three-wheeled iguanas" and the like. The psychology of imaginary behaving systems has since become the centerpiece of Braitenberg (1984). His vehicles, though highly idealized, illustrate the advantages of "synthetic psychology": A system we *build* is much easier to understand than one we *find*. I suggest that *Rana computatrix* be located in the phylum of these *Gedanken* beasts, the most advanced of its kind yet to evolve.

Accordingly, whether *R. computatrix* is a good model of living Ranidae or Bufonidae is a secondary issue, because the model is *itself* of cognitive interest. Especially in its forthcoming incarnations (sections 5 and 6), *R. computatrix* will be exhibiting simulated behavior in its simulated environment that will invite cognitive description and analysis. In particular, models like Arbib's might be valuable tools for addressing foundational issues in cognitive science and philosophy. High on the list is the understanding of *representation*. Arbib writes of internal models of the world (section 2), assemblages of instantiated schemas that correspond to "domains of interaction." How does this picture account for features of the *specificity* of representation? Do the instantiated schemas in *R. computatrix*'s brain represent *flies, small distal objects*, or *small patches of retinal stimulation?* Similarly, under what conditions will we say that *R. computatrix*'s model of the world is *false?* These are central issues for cognitive science, and complete models like Arbib's can provide concrete guidance in their resolution. (For a more extended discussion, see Lloyd 1987.) Progress on these issues will contribute ultimately to specifying the continuities and differentiae of animal and human cognition.

In sum, Arbib's approach is incremental in several senses: He assembles model systems part by incremental part; both top-down and bottom-up data incrementally determine the models; new data lead to incremental adjustments; the models – of both living systems and *Gedanken* beasts – fall in an incremental continuum of models of interest to cognitive science. Each progressive enlargement is a contribution, and the whole cognitive science that results is greater than the sum of these incremental parts.

What is the schema for a schema?

Alan K. Mackworth
Department of Computer Science, University of British Columbia, Vancouver, B.C., Canada V6T 1W5

Arbib's main thesis is that schema theory can serve as an intermediate level of functional analysis. His subsidiary thesis is that it should so serve because it thereby provides a common language for cognitive science. The argument that it *can* is by demonstration. The argument that it *should* is by exhortation.

Since the modern introduction of the concept of schema by Head (Head & Holmes 1911; Oldfield & Zangwill 1942) the concept itself has repeatedly demonstrated an ingenious talent for metamorphosis. A schema has variously been identified with a map, a record, a pattern, a format, a plan, a conservation law (and a conversation law), a program, a data structure, a coroutine, a frame, a script, a unit, and an agent. Each of these concepts has, in turn, considerable variability and ambiguity. The schema for a schema is extraordinarily plastic. "An organism has somehow to acquire the capacity to turn around upon its own schemata and to construct them afresh" (Bartlett 1932, p. 206). Our schema for a schema certainly demonstrates that capacity.

Unfortunately, if we allow this unbridled plasticity, the main thesis becomes impossible to falsify, making it a slogan rather than a scientific hypothesis. Arbib is aware of this trap but, still, his concept of schema must be sharpened up and the arguments strengthened. The argument for the main thesis hinges on establishing a single, encompassing but precise schema for "schema." The argument for the subsidiary thesis hinges on the same point. Schema theory risks suffering the fate of the Tower of Babel. It cannot, in its current state, serve as the common language of cognitive science.

Arbib provides some functional criteria for schemas and four structural definitions that, although precise, are far from coextensive. They can mostly be characterized as defining a schema as a program in a serial programming language with some mechanism for parallel execution of subprograms.

So our primary goal should be to evolve a coherent theory of schema systems that, as Arbib emphasizes, must serve the needs of functional and structural analyses. His target article is clearly a major step in that evolutionary process. One weakness, however, is in his following the prescription of identifying necessary functional attributes of a schema system. The four criteria given are somewhat vague and nonspecific, which leads to the range of structures and languages proposed as realizations of schema theory.

One source of the problem can be identified. In rejecting Marr's notion of the "computational level" as an independent level of analysis Arbib ignores a rich source of functional constraints. Although the term "computational level" is a very poor choice of words, the idea itself is crucial. Let us call it, instead, the task level.

Analysis at the task level allows us to determine the natural constraints that *may* be exploited by a system and the structure of the equivalence classes of solutions to the perceptual or motor task. We can determine whether a solution is overconstrained or underconstrained by the various potential information sources. If it is underconstrained (that is, the solution equivalence classes have more than one member), then additional constraints or preferences can be identified that may be imposed to select a particular solution (Mackworth 1983).

Task analysis also allows us to determine the inherent task complexity – the degree of difficulty of the problem itself with respect to a family of computational architectures – using measures of time, space, number of gates, degree of parallelism, and so on. These tools put necessary limits on any algorithm's performance as well as allow us to analyze the actual performance of a particular representation and algorithm. There are intriguing new results from theoretical computer science to be exploited. For example, the appeal to massive parallelism is a standard escape hatch these days, but complexity analysis of tasks reveals that some stubborn problems may be inherently serial in nature, not to be dissolved by the magic of massive parallelism. Just as Minsky and Papert (1969) showed that connectedness is an inherently serial problem within the framework of an earlier, more restricted, parallel scheme, the Perceptron, we should be aware of the existence of this class of "intrinsically" serial problems and the implications for schema theory. For example, establishing a form of local coherence, arc consistency, between a set of declarative schema instances is, surprisingly, such a problem (Mackworth 1987).

Functional analysis at the task level can lead to the identification of adequacy criteria that *any* representation scheme must necessarily satisfy. Eleven descriptive adequacy criteria and five procedural adequacy criteria are identified and discussed in Mackworth (1987). These can serve as functional criteria for schema theory (Havens & Mackworth 1983).

Even given highly constraining functional criteria, there are, of course, an unbounded number of schema theories that could satisfy them. Arbib emphasizes the need to consider also the constraints arising from the neural architecture. Are the four schema theories he proposes so constrained? They assume, for the most part, "von Neumann" architecture, including subroutines called by name with numerical parameters, sequential control with program counters, and minor modifications for parallelism, including the FORALL mechanism and coroutines. Are they appropriate languages to be implemented on the cognitive architecture, given what we know of it? There is a chasm between the schema level and the neural level. Bridging it is a major task in the research program Arbib lays out. The relationship established in Section 5, between a schema theory of prey and obstacle detection in the toad and its neural implementation, is a rickety bridge but it is one of the few we have.

In summary, we must apply the functional and structural approach to the analysis of schema theory itself. It is critical that we establish the tightest possible functional criteria for schema theory and derive a crisp structural theory: the schema for a schema. The claim that schema theory is a necessary component of the analysis of behavior thereby becomes an empirically testable hypothesis.

ACKNOWLEDGMENTS
The author is supported by a fellowship from the Canadian Institute for Advanced Research and grants from the Natural Sciences and Engineering Research Council of Canada.

Eye of toad, and toe of frog?

John C. Marshall
Neuropsychology Unit, Neuroscience Group, The Radcliffe Infirmary, Oxford OX2 6HE, England

I respond to Michael Arbib's target article as one whose concern with frogs is somewhat limited (limited in fact to recipes for cooking their legs), but whose interest in visually guided behavior is considerable. The reason for this latter concern is, of course, that, as a hospital-based neuropsychologist, I see brain-damaged patients with a wide range of (acquired) disorders of visuospatial cognition and praxis.

To mention but a few illustrative examples: In patients with bilateral parietal lobe damage, there is often "global spatial disorientation" (Kase et al. 1977). Despite normal language and intellect, these patients have randomly wandering eye movements, they fail to track slow-moving objects with their eyes, and are unable to touch or grasp objects reliably with either hand, overshooting, undershooting, or veering to left or right. Such patients may also be unable to perceive visually more than one object at a time (Berti et al. 1986), and may misorient to auditory stimulation. In addition, there can be associated disorder of topographical learning and memory (patients failing to find their way from corridor to ward), and, upon being guided to the bed, such patients may be quite unable to orient themselves into a reasonable position for lying down (despite adequate object recognition).

Some of these deficits (and many others) can be seen in relatively "pure," isolated form. Thus Ratcliff and Davies-Jones (1972) demonstrated that defective localization of stimuli in peripheral vision can be dissociated from generalized visual disorientation; deficits in visually guided maze learning, in which maze and patient are in a fixed relationship, can dissociate from locomotor map-following, where the patient's orientation relative to the environment changes (Ratcliff & Newcombe 1973). Right posterior damage can provoke a particularly salient inability to mentally rotate objects in space (Ratcliff 1979). And topographical memory loss, consequent upon right medial temporo-parietal lesions (Landis et al. 1986), can be found in the context of many relatively well-preserved perceptual and spatial skills (Whiteley & Warrington 1978).

A severe inability to perceive two objects simultaneously ("simultanagnosia") can coexist with good ability to recognize single objects after bilateral occipito-parietal injury (Luria 1959); left-neglect, where the patient with right posterior damage behaves as if the space contralateral to the lesion does not exist, can present in a variety of forms; the underlying impairment may be primarily attentional (Baynes et al. 1986), motor (Joanette et al. 1986), or representational (Bisiach & Luzzatti 1978). The condition further fractionates into disorders of extrapersonal and personal space (Bisiach et al. 1986). With respect to personal space, very selective autotopagnosias (difficulty with pointing to body parts on command) have been reported after left parietal damage (Ogden 1985) in nonaphasic patients. The role of the left hemisphere in some visuocognitive skills is also brought out by the fact that severe impairment in the detection of "hidden (geometric) figures" is preferentially associated with aphasia-provoking lesions (Orgass et al. 1972; Teuber & Weinstein 1956).

And finally one might mention that in "speakers" of American Sign Language (Poizner et al. 1987), spatial syntax in the service of linguistic communication can be firmly dissociated from the perception and manipulation of objects in extrapersonal space; sign-language aphasias are consequent upon left-hemisphere damage (as are spoken-language disorders), whereas right-hemisphere damage results in nonlinguistic impairments of spatial topography.

These, then, are just a few of the visuocognitive and visuomotor deficits concerning which neuropsychologists would be grateful for any theoretical (or perhaps even computational) insights that anyone can offer (for more extensive review see De Renzi 1982; Miller 1986; Young 1983).

Will toads, frogs, and their artificial analogue, *Rana computatrix,* provide that help? The overall structure of Marr's theory of vision (Marr 1982) has already provided us with new and helpful ways of looking at the varieties of visual agnosia (Humphreys & Riddoch 1987; Ratcliff & Newcombe 1982). Will *Rana computatrix* play a similar heuristic role in elucidating other disorders of visual cognition? Or is the prey-catching, predator- or barrier-avoiding behavior of the frog based upon mechanisms that are too special-purpose (and too closely linked to the specific anatomy and behavioral adaptations of the frog) to be of use to me and the patients I see? I do not know. But I do wish that more members of the simulation community would take a look at the kinds of data provided by disorders of human cognition.

As I said, we need all the help we can get. If we are ever to provide serious remediation for patients with these often devastating disorders (either by "retraining" a damaged system, or by finding alternative ways around the impairment), we shall require models vastly more sophisticated than those that I, at least, currently work with. I would like to think that *Rana computatrix* is a step in the right direction; but frogs will not excite me in nonculinary ways until I can see the relationships between this kind of modeling and the problems with which our patients present. Is global stereopsis in the frog linked with the perception of anomalous (subjective) contours? Both normal and pathological data suggest that the functions are intimately related in man (Hamsher 1978). When Arbib stresses "implementational details" and notes that different animals "may make different use of visual cues," I begin to worry whether *Rana computatrix* does it our way.

Recent physiological findings on the neuronal circuit of the frog's optic tectum

Nobuyoshi Matsumoto

Department of Biophysical Engineering, Faculty of Engineering Science, Osaka University, Toyonaka, Osaka 560, Japan

Arbib introduces neuronal circuit models for the frog, and discusses how visually guided behaviors are generated by them. The interpretation of behaviors by means of neural microcircuits is of great interest to physiologists. Unfortunately, data are not sufficient for discussing the behaviors precisely. Arbib depends mainly on the tectal column circuit proposed by Székely and Lázár (1976) in describing *Rana computatrix.* It should be remembered that the tectal column structure was originally proposed mainly on the basis of anatomical studies; no detailed physiological information had been available at the time. Nevertheless, the circuit model has some significant implications not only for anatomy but also for physiology. I will present some supplemental physiological data relating to the neuronal circuit of the frog's optic tectum that might help strengthen the model.

First of all, there is a point I do not clearly understand in the modeling. A glomerulus shown in Arbib's Figure 2(d) is connected directly to two dendritic endings that belong to different neurons. Anatomically, they should make contact to form synapses. I am afraid that this kind of simplification might lead to some errors in evaluating the function of the model, especially when discussing the time sequence of excitatory and inhibitory potentials.

As Székely and Lázár have pointed out, there are at least two restrictions on the neuronal circuit diagram of the tectal column. First, large ganglionic neurons are not involved in the circuit. Our recent results (Antal et al. 1986; Matsumoto, Schwippert & Ewert 1986), however, have shown that large ganglionic neurons play an important role in information processing in the optic tectum. Unfortunately for the model, large ganglionic neurons showed much more variety in morphology and physiology than we had expected.

Second, the contribution of class 3 retinal fibers to the tectal column is not considered. Physiologically, an interaction between the myelinated and unmyelinated fiber system has been observed in many of the tectal neurons recorded (Matsumoto & Bando 1978; 1980). One of the possible interactive circuits is shown in Figure 16(b); this model was used to explain the discrimination between wormlike and antiwormlike stimulus patterns. We identified the prey-selective efferent neuron as the pyramidal neuron and found that they receive a strong inhibitory input from class 3 retinal fibers (Antal et al. 1986; Matsumoto & Antal 1984; Matsumoto, Schwippert & Ewert 1986). There is an indication, as shown in Figures 7 and 16(b), that thalamic neurons are involved in this inhibitory circuit (Ewert 1984). However, we have no further information about whether the thalamic neurons directly inhibit prey-selective (pyramidal) neurons or whether they inhibit through tectal interneurons.

Arbib has shown facilitation of prey-catching behavior using his model. One interesting phenomenon that might be related with facilitation should be mentioned. Suppose we electrically stimulate the optic nerve. The large pear-shaped neuron in the tectal column produces a monosynaptic and a disynaptic EPSP (excitatory postsynaptic potential) followed by an IPSP (inhibitory postsynaptic potential). In our earlier experiment (Matsumoto & Bando 1980), exactly the same response could be recorded from some of the large ganglionic neurons. Presynaptic fibers were found to be myelinated rather than unmyelinated ones. Since the conduction velocity of unmeylinated fibers is too slow to make a reliable analysis, we do not know whether or not the same effect can occur in the unmyelinated fiber system. It is easily understood that reverberatory responses could be triggered if the inhibitory potentials were eliminated. When large ganglionic neurons were visually stimulated, some of them showed reverberatory responses (Matsumoto, Schwippert & Ewert 1986). If the stimulus suddenly stops in the receptive field, the amplitudes of EPSPs are greatly enhanced, probably because of focal stimulation of the excitatory receptive field. This effect was greatest for a wormlike stimulus, which produces the smallest inhibitory effect on the neuron. We do not know, however, whether the reverberatory response seen in large ganglionic neurons is dependent on the positive feedback circuit in the tectal column.

As mentioned earlier, we physiologists have not collected enough data for a detailed analysis of the function of the optic tectum of the frog. To solve this problem we are now investigating synaptic connections between retinal fibers and tectal neurons by recording simultaneously from presynaptic (extracellular) and postsynaptic (intracellular) elements (Matsumoto, Nagano & Li 1986). We soon hope to be able to provide new findings to improve the model of the neuronal circuit of the frog's optic tectum.

The computing frog

G. Székely
Department of Anatomy, University Medical School, H-4012 Debrecen, Hungary

Although we are still very far off from understanding the neural mechanisms underlying visually guided behavior of the frog, I find Arbib's endeavor a fortunate choice for modeling the visual part of the brain, which has been in the forefront of interest for the last few decades. The basic idea is to describe behavioral phenomena in terms of interacting schemas as units of computational analysis and then to implement the schemas, perceptual and motor, in realistic models of neural circuitry. I find this approach a very plausible and promising one, and the strong commitment of the author to remaining in continual contact with the biological material should receive the greatest recognition from experimenting scientists. The target article raises several questions, of which I would like to discuss two in my commentary.

1. How realistic is a "realistic" network? This question necessarily comes to a neuroanatomist's mind whenever he tries to extract some kind of "biologically meaningful" neuronal organization from a body of crude experimental findings. To make it quite clear, the structure of the nervous system is so complex that the experimentalist is compelled to make, deliberately or unconsciously, conceptual models in order to describe and interpret his results. In such conceptual models the invariants of the universal are abstracted, simplified, and arbitrarily rearranged to fit the preformed concepts of the experimentalist. This applies just as much to the simple reflex arc and the several ingenious "conceptual models" of Ramón y Cajal (1911) as to the columnar organization of the optic tectum that I and Lázár (1976) described.

The reason I so explicitly point to "conceptual modeling" – to the probable disagreement of many of my experimenting colleagues – is that some aspects of the perceptual schema domain are modeled back to the assumed "reality" with the help of my tectal column concept. Although I tremendously enjoyed the lovely computer experiments of Cervantes-Perez (1985) on these columns, I could not help uneasily feeling how much he had to change and twist the (otherwise unknown) parameters of the elements of the column before he obtained the sorts of responses from the pyramidal neuron that Jörg-Peter Ewert (see his target article in this issue) would have expected of his T5(2) neuron, (which, being a "worm detector, " is also a conceptual model). The same applies to the likewise beautiful computer experiments on the 8×8 array of tectal columns (Cervantes-Perez et al. 1985), which, however elegant, are disappointingly trivial in showing the retino-pretectal-tectal interaction exactly in the same way as Ewert would unambiguously describe in three sentences.

In both computer experiments success consisted of the best possible simulation of the animal's behavior. This is exactly what one would expect of good modeling if the premises – the tectal column, the physiological classification, and so on – represented real invariants of the world. The tectal column, for example, is one of the many possibilities among how the neuronal interconnections may be organized in the optic tectum. In my opinion, a computer scientist should not stop at this concept; he should perform the role of a computer *experimentalist*, using his machine ability to test such possibilities. He should not take all biological data (the conceptual models) at face value but, with the freedom of the experimentalist, try to see whether other neuronal interactions would yield better, or worse, results in terms of some arbitrary criteria. He may then call the attention of the biological experimentalist to prior faulty conclusions, to overlooked data, to the need to reinterpret existing data, to some "emergent" properties that remained unnoticed in biological experiments; and a number of other interactions may evolve. The computer scientist just waiting for biological data is a passive partner to the biologist, as opposed to the *computer experimentalist*, who not only checks the biologist's data and interpretations but provides him with ideas and inspirations for new experiments.

The biologist facing the problem posed by thousands and thousands of interacting elements acutely feels the burden of uncertainty in his approach, and needs the contribution of the computer scientist just as he needs solid biological data. It seems to me that *Rana computatrix* would be an excellent tool in such an active partnership. Even in its present state, *R. computatrix* calls for more detailed correlative morphological investigations, which have only just begun with two recent publications (Antal et al. 1986; Matsumoto, Schwippert & Ewert 1986). Serial EM (electron microscopic) investigations we have just initiated in this laboratory are badly needed to verify and explore neuronal interconnections; above all immunohistochemistry combined with other neuronal tracing techniques would yield important data for *R. computatrix.* I am sure physiologists have other lists of data. Conversely, a number of experiments could be performed with *R. computatrix* (e.g., giving inhibitory character to different neuron types, incorporating modulatory effects of various peptides, etc.).

2. The investigation of visually guided behavior is almost always focused on the optic tectum, which is certainly an important visual center but not the only one in the frog. In addition to the pretectal structure, there are two more areas receiving optic fibers in the anterior thalamus and one in the mesencephalic tegmentum. Except for this latter center, all other areas receive a precise retinotopic projection of the visual field and are reciprocally interconnected in a remarkably organized fashion (Lázár 1984; Székely 1971). Highly organized reciprocal interconnections exist between the tectum and the isthmic nucleus, and the role of this latter nucleus is only recently appreciated in visually guided behavior (see Arbib's target article). These many centers with highly organized interconnections make one wonder about their role in the control of visually guided behavior. As an experimenting biologist (though perhaps not wholly aware of the difficulties a computer simulation may create) I would like to see the incremental evolution of *R. computatrix* in this direction as well.

Schema theory: A broadening viewpoint

Tang Yi Qun
Department of Computer Science, Wuhan University, Wuhan, Hubei, The People's Republic of China

Research on the brain, behavior, and intelligence needs intercommunication among several disciplines; but cognitive psychology and artificial intelligence do simulation based mainly on functional levels, whereas neurophysiology and anatomy are interested in neurons and neural circuitry. There is a gap between them. Arbib has advanced schemas as a bridge to link them: Functional modeling could gain more insight into structure from neurons and neural circuitry through schemas; on the other hand, experiments on neurons could receive some guidance from more global viewpoints. To play the role of a bridge, schemas should be compatible with both functional modeling and experimental data on neurons and their circuitry – that is, they must not only be able to represent human thought and behavior but must also map into some areas of the brain in a one-to-many or many-to-many fashion.

If schemas are considered as units representing and interacting with the world, then it may be possible to represent human thought and behavior with an adaptive schema network. Fortunately, certain functions and their corresponding subnetworks are clearly hierarchical. Schemas located at the same level would interact by competition and cooperation, whereas schemas at different levels would act by maintaining the proper sequence. For example, motor control – from the motor cortex to the cerebellum, brain stem, spinal cord, and muscle servo – is clearly hierarchical. [See Stein: "What Muscle Variable(s) Does the Nervous System Control in Limb Movements?" *BBS* 5(4) 1982.] Visual information processing in the frog and the toad is another example. On the other hand, in each level of the hierarchy there are many concurrent activities, for instance, the information processing in retina. The brain accordingly processes and integrates information in both sequential and parallel distributed ways. Arbib's operations on schemas reflect this feature, such as compounding (several schemas compounded are a more powerful new schema), generalization, instantiation, and activation (data-driven or goal-driven).

I think it is important to find more methods for representing schemas. Semantic networks are successfully used to represent structured, well-learned knowledge in artificial intelligence, but they are suitable only for the analytic mode of thinking that breaks the object up into parts and then combines them. Dividing the whole into parts may lose much information. There is another mode ("imagistic thinking") whose function is to process the object as a whole (Qian 1983). This helps to abstract similarities among objects, as used in analogies and logical inference. What schema representation will facilitate such analogies and abstraction in this mode of thinking? It is likely to be quite different from those that are suitable to the analytic mode.

The ideal of using the schema concept as a bridge is in an early stage of development and perhaps somewhat premature, but worthy of further exploration.

Schemas: Not yet an interlingua for the brain sciences

John K. Tsotsos

Department of Computer Science and The Canadian Institute for Advanced Research, University of Toronto, Toronto, Ontario, Canada M5S 1A4

Arbib argues for the utility of schemas as intermediate representations between neural-level analysis and overall brain behavior. He presents examples on how to move from schemas to testable models of neural circuitry, proposing that schemas can provide a meeting ground for the behavioral and brain sciences. Unfortunately, there is no reason to believe that there is only one requisite level of intermediate representation. I agree that there is a need for a common language of discourse, and for stating theories. Only with a common formalism can theories from one discipline contribute to development in another. In other sciences – for example, physics – the language of mathematics forms this common formalism. However, schemas, as Arbib uses them, do not yet provide this interlingua for the brain sciences.

The key problem with Arbib's schemas (as well as with Minsky's frames [Minsky 1975], Schank's scripts [Schank & Abelson 1977], and the other versions of packaged knowledge) is that a syntax and semantics for the formalism are not provided. Arbib only provides a set of four constraints on the characteristics of schemas. It is difficult to argue against the utility of representation or of knowledge packages in general; they have so many advantages and these have appeared in many papers (a partial summary appears in Tsotsos 1984). There is a need, however, to introduce the concept to other domains because, although it is perhaps well accepted in the computer science community, it may not be so well accepted in other disciplines. The designer of a representation must determine how appropriate the representation is for the problem domain, and must give it a syntax and formal semantics so that others may adopt it as well. A common formalism is useful only if other researchers use it, if one can compare different representations of the same problem solution, and if one can distinguish correct from incorrect representations. Throughout most of the target article, Arbib uses schemas in an informal manner; a great many questions arise about schemas in general and about the representation of specific constructs. The general questions are:

How are schemas mapped onto an implementation for testing purposes? If they are part of a programming language, is there a compiler or interpreter for the language?

How can schemas proposed for the same behavior be compared without simulation?

What class of behaviors does a given schema represent? How can one tell whether that class of behavior is the desired one, and contains no other unwanted behaviors?

How can one tell whether a given schema formalism is sufficient for all the behaviors of interest, and only those?

The specific questions are:

Is a schema a procedural or declarative entity? Does this matter?

Are the arrows between units in a schema control lines, data lines, or do they represent some other relationship?

What are the primitives of schemas? Is it possible to define a set of primitives so that all behaviors of interest can be represented using combinations of these primitives?

In motor tasks, temporal issues are of particular importance. How is time represented in the schemas?

How are motion qualities, such as "grasp gently" or "grasp aggressively," represented?

What are the subtleties of vision that the language of Arbib's Appendix B cannot capture, in addition to not handling control concepts? Are they not critical?

It is claimed that the "FORALL" construct is novel; it appears previously at least in my own work, and definitely in my 1984 contribution to Arbib and Hanson's edited volume (Tsotsos 1987). In fact, the work of Weymouth described in Arbib's Appendix A is very similar to the representation and control framework for vision I had put forward (Tsotsos 1980) and further refined (Tsotsos 1985; 1987). In fairness, however, my own work does not provide answers to most of the questions I have raised here (and no one else's does either, I believe).

The problems with schemalike and semantic net representations were first noted in seminal papers by Hayes (1974) and by Woods (1975). Further analyses are found in papers by Brachman and by Hayes (Brachman 1983; Hayes 1979). These papers launched a very different emphasis in the knowledge representation research community: that of providing representational formalisms with a logical semantics. (See Mylopoulos and Levesque [1984] and Levesque [1987] for overviews of knowledge representation research and Brachman and Levesque [1985] for a collection of key papers in the field.)

The need for a common language of discourse among the artificial intelligence, cognitive psychology, and neuroscience communities is critical, particularly with the current burgeoning interest in many centers in interdisciplinary research. In this regard, Arbib is completely right, and I regard his proposal of schemas as a very reasonable starting point. The questions I have raised are not so much criticisms of Arbib's work but indications that much more work remains to be done. It is appropriate to examine the success of mathematics as a common language for the physical sciences. Mathematics has an agreed-upon formal semantics. A family of compatible representational tools is provided, each suitable and appropriate for a different problem task (arithmetic, algebra, geometry, topology, probability, and statistics, calculus, analysis, set and number theory, logic, combinatorics, and others). I feel it is a mistake to expect that a single representational formalism can be found that will be sufficient and appropriate for all levels of analysis of the brain. I believe that a family of representations is required in much the same way as a family of representations and tools has been found to be appropriate for the physical sciences. The members of this family must be compatible with one another in the same way that the different languages of mathematics are compatible. Schemas or their derivatives are certain to play a role at some level of representation.

Schemata and representational constraints

Cees van Leeuwen
Department of Experimental Psychology, University of Nijmegen, 6525 HR Nijmegen, The Netherlands

The target article gives me the impression that the author is telling the wrong story at the right time. Arbib acknowledges a point not generally accepted by neurologists – that a functional explanation of a neurological system is necessary in order to capture the relevant aspects of its structure. He therefore accepts a cognitivist type of description as intermediate between the behavioral level and the neurological one. Cognitivism identifies functional description with formal description. The formal descriptions involve computations on representations. Since "the eventual goal, of course, is for functional and structural analyses to be rendered congruent," it comes as no surprise that Arbib seeks to combine neural circuitry into more complex ensembles called "schemata." However, he does not wish to accept the computational level of explanation as a separate level, because "different animals (or different subsystems of a given animal) may make different use of visual cues that cannot be discovered until 'implementational details' are taken into account." If the role for computational notions is merely a heuristic one, they will ultimately be eliminated from the theory. It seems, however, that they are meant to stay, since "For many behaviors, analysis at the level of single neurons may be superfluous." Arbib therefore seems willing to have the best of both worlds, without even mentioning the conceptual problems involved in identifying the computational level of explanation with a certain type of neurological hardware.

But let us dream of such an account with all these problems solved. Basic steps in the computations are carried out by the underlying neurological system. We assume that its "implementational details" are simply "facts" about the neurological hardware. We assume that these facts indeed constrain our elementary computational steps. Then why do schemata need to play the role of the elementary computational processes provided by the neurological hardware? They seem far from elementary processes.

If this is the point Arbib wishes to make, my critique boils down to the following: The question of which processes are carried out depends on the question of which representations can be stored and retrieved. Constraints on representations are therefore at least as fundamental as those on processes. An interesting theoretical possibility lies in imposing neurological constraints that have consequences for representations. Satisfaction of these constraints would be a prerequisite for the ascription of meaning to representations in a cognitive explanation. This means that the constraints themselves should be formulated independently of any semantic domain. Any domain-dependent formulation of the representational properties is beside the point. Schemata are, alas, domain-dependent entities, and can therefore not be expected to serve as a basis for understanding the neurological constraints on computational theories.

Schema theory: A new approach?

W. von Seelen
Arbeitsgruppe III (Biophysik), Institut für Zoologie, Johannes Gutenberg Universität, D-6500 Mainz, Federal Republic of Germany

As his books have demonstrated, Arbib is familiar with theories no matter how complicated they are. One therefore wonders what impels him to construct models – for example, models of simple toads—that do simple things in simple experiments. Many physicists have avoided doing just this and have immigrated into brain research (e.g., from thermodynamics) with the aid of complicated equations. On the basis of similarities between "physical" and "neuronal" structures, they have then attempted to explain or at least to imitate brains. These attempts have been "successful": They have led many scientists to think about information processing in terms of a kind of neuronal "architecture" and thereby to solve problems such as pattern recognition or optimization algorithms that are useful in technical applications. However, such work has repeatedly arrived at the conclusion that almost nothing about the brain can be explained in this way. Given the conditions of the real world, the equations in physics are approximations that are frequently inadequate when boundary conditions alter. Physics is certainly essential for the understanding of biology, but in its conceptual methods rather than in uncritical use of its equations. The result is that there are many formulas in brain research today, but few concepts in which theories can take root.

In my opinion, Arbib's target article presents a concept that could be seminal. It is justified eloquently and illustrated with many examples. In order to comment on it – insofar as I have understood it – I would like to "distort" it a little as seen from my own point of view: To describe brain functions, a certain level of abstraction is necessary. Arbib chooses basic schemata, which have to be constructible in neuronal architecture that is technically measurable. As a help in finding and defining basic schemata, I think they should be interpreted as "basic situations." These consist of sensory or motor cues or both. To perceive and handle these situations – for example, prey-detecting and prey-catching – is the task of the animal. By combining basic situations, more complex behavior can be generated. The spatiotemporal behavior of layered feedback structures can be used to encode and combine such representations, as Arbib shows for the stereo problem. This approach assumes that brains are not universal computers but rather systems that must solve concrete problems in a specific environment, that must be capable of evolution, and that can be constructed reliably under the constraints of ontogeny. The development of *one* new schema allows a large number of combinations with already existing ones, and the requisite *hierarchies* are easy to construct. Compared with the algorithmic approach, this concept, based on the structure of the relevant task and neuronal networks, has the advantage of using experiments that can be currently implemented as well as computer simulated.

This approach of course also involves problems not solved in the target article:

1. How can basic schemata or basic situations be defined and systematically detected?
2. What strategy is used to couple the basic schemata when a target function has to be defined?
3. Is there a general and workable data format for such systems?
4. Is self-organization possible?

I think it is worth proceeding in the direction Arbib has outlined. The path is difficult, but it is the one physics has used since the beginning of experimentation.

The computer simulations tend to conceal a conceptual detail I think should be reexamined: the columnar structure of neuronal networks. This has not been detected neuroanatomically and may be misleading functionally. Neurophysiological recordings suggest that there are columns in a layered structure with retinotopic mapping; but because there is feedback within and between layers, such columns would be functionally irrelevant. Since the different layers have separate outputs (cortex), their functions are also different; and columns are not adequate as an implementation of an operator for a picture point, for example. A more viable hypothesis may be a neuronal layer as a basic system complemented – if necessary – by discrete, geometrically arranged grids to which input is directed or where specific internal processing takes place (e.g., "blobs" in the visual cortex).

In view of the almost uniform structure of large brain areas (e.g., the neocortex of mammals) it is possible to get by with a few kinds of networks. However, these have to be analyzed and implemented as completely as possible. So, as far as the toad is concerned, one should investigate whether mapping the retina onto the tectum using coordinate transformations – as occurs in mammals – would not considerably simplify processing, especially with moving stimuli.

Schemas and bridging gaps in the behavioral and brain sciences

Johan P. Wagemans
Laboratory of Experimental Psychology, University of Leuven, B-3000 Leuven, Belgium

Arbib's target article is an extremely valuable effort in bridging gaps among disciplines, theories, and topics. The elaboration of this claim will be the major focus of this commentary. Some shortcomings or difficulties are also pointed out, but compared with the extremely high value of the major part of this work, the latter are only minor quibbles.

1. Major achievements: Bridging gaps. In the behavioral and brain sciences cross-disciplinarity is often argued for but seldom achieved. Indeed the gap between the cognitive sciences studying the higher-order processes of human minds and the neurosciences studying the neural tools embodying these processes is large and deep. In his research on the mechanisms of visually guided behavior, Arbib lays the groundwork for a connection not only among disciplines, but also between two often divergent strategies of doing science (theory building versus experimental testing), between metatheoretical approaches (ecological versus constructivist), and between research domains (perception and action).

Concerning the (lack of) relation between the cognitive and neurosciences, Arbib clearly states that the higher-order behavioral and cognitive capacities and processes cannot be reduced to low-level neural tissues and firings in one single step. Some intermediate format is needed in gradually bridging the gap between the two extremes. Arbib argues for the notion of schema as a useful candidate in doing this job. Considering the fruitfulness of this conjecture in the empirical achievements of his work, Arbib may well be right. In using incremental modeling, as he calls it, Arbib also demonstrates that constructing theories and doing experimental research are not as divergent as is often the case. Although every scientist would claim that the ideal strategy of doing science is certainly the combination of the two, those who are actually good at both are very rare.

Arbib is not unique in stressing the need for a level of analysis mediating between the strictly cognitive and neurophysiological ones. Others have made a comparable plea for a third level of explanation in the behavioral and brain sciences: For example, Newell (1980; 1982) has argued for a symbolic level between the physical or neurophysiological level and the knowledge level. In a similar vain, Pylyshyn's (1980; 1984) view on cognitive science is centered on a distinction of (1) basic mechanisms, (2) a syntactic or computational level, and (3) a semantic or representational level, all three of which are essential in capturing generalizations that cannot be made on another level. The difference between Arbib's and other approaches lies in the details offered in the research on all three levels and in the categorizing of the levels. Whereas a schema would be considered an instance of the highest level (knowledge, semantic or representational) in the mainstream of cognitive science, Arbib sees it as an intermediate-level construct and, indeed, as he specifies it for the perceptual and behavioral domain, he might well be right: The content of the schema is determined by its formal computations and as such it is not semantic in the strict sense.

The notion of a schema is a relevant one in considering the difference between direct and indirect theories of perception (and broadened to the total domain of the behavioral and brain sciences, the divergence between ecological and constructivist or cognitive approaches). Gibson (e.g., 1979) would have called schemas, like any other intervening processes or structures, superfluous in explaining perception or other psychological processes. Nevertheless, the way Arbib uses the term – namely, as a process whereby "the system determines whether a given domain of interaction is present in the environment" – is much more compatible with the ecological approach than would be expected of the schema notion as such. The example of prey selection or obstacle avoidance could easily fit the ecological framework of detecting the edibility or pounce-on-ability of the stimulus patterns by the resonance of the perceptual system with the invariants specifying those affordances. The huge step forward of Arbib's approach, in comparison with Gibson's, is that Arbib is much more explicit about the way the tuning actually occurs. In using his incremental style of modeling, he is able to specify even the exact neurophysiological structures and processes embodying the schemas involved in the perceptual pickup and its behavioral consequences.

The focus of Arbib's target article, visually guided behavior, represents a serious and detailed effort to bridge the gap between perception and action. This problem was a critical one in the early days (cf. Tolman's rats buried deep in thought; see Guthrie 1935, 1960), as well as very recently (cf. the latest research efforts in robot vision, e.g., Brady & Paul 1984; Horn 1985). Attempts to solve it accordingly deserve praise when the theoretical models are as detailed and empirically grounded as Arbib's research on visually guided prey-catching and obstacle avoidance in the frog and the toad.

2. Minor shortcomings and difficulties. As I warned in the introduction, my remarks are only minor and disparate in nature and do not take anything away from Arbib's major positive contribution to the progress of our sciences.

Arbib's account of Marr's (1982) computational investigation of visual information processing is too negative and based on a serious misunderstanding of the specificity of the computational level of analysis, Arbib argues that Marr is wrong in regarding this computational level as independent. He states that "one cannot give an a priori analysis of depth perception because different animals . . . may make different use of visual cues that cannot be discovered until 'implementational details' . . . are taken into account" (sect. 1, para. 4). Now, Marr's claim is exactly that the same computational task (e.g., of recovering depth from 2D images) can be done by several algorithms (e.g., shape from shading, structure from motion, depth from stereo-images) that each have distinct implementations. So the finding that different organisms use other "depth cues" is no refutation of Marr's approach but a clear confirmation. Arbib's view that Marr's ideas add nothing to the virtues of classic functional accounts is equally mistaken. The level of mathematical rigor and computational detail attained by Marr and his group is unique and was never before achieved.

Arbib is not right in characterizing modules as structural units (sections 1 and 6). As stressed by Fodor (1983; 1985), and as emerges from Marr's (e.g., 1982) work, a module is, on the contrary, a strictly functional unit. That is precisely the difference between the relatively new concept of module, defined by computational criteria, and the rather old notion of layers, structurally defined.

In addition to these failures in Arbib's target article, perhaps he could specify in his response some further details concerning several insufficiently worked-out statements. First, in the section on prey selection (3.2), Arbib discusses Didday's (1976) model hypothesizing "sameness" and "newness" cells, suggesting that further experiments need to be done to work out this conjecture. Our question is whether it is not a logical mistake to suppose the existence of two different cells to detect the two values of one single Boolean variable, namely, new or not. Second, in the section on prey acquisition (3.4), Arbib attributes two discrete values, d+ and d−, to the continuous distance parameter. Must these values be specified a posteriori, or is it possible to connect them with a kind of schema for the snapping range of the frog's tongue? Finally, Arbib stresses the need for a distributed model (sect. 4.2, para. 1) and claims that cooperative computation is a general principle of brain operation (sect. 5, para. 1), but compelling theoretical or empirical arguments supporting these statements are seldom offered and should therefore be provided in the Author's Response.

Grasping schemas is (are) difficult

H.T.A. Whiting
Department of Psychology, Faculty of Human Movement Sciences, The Free University, Amsterdam, The Netherlands

Whereas both Thelen (1986) and Arbib (in his target article) would seem to agree in interpreting behavioral outcomes at any point in time as products of dynamic, relational, multilevel interactions among system properties, Thelen's concern has been not only with behavior in the real (movement) time domain but also over the developmental time span. It is perhaps this double perspective that has led her to interpret the complex ordering and co-ordering observable in both moving and developing organisms – conceived as cooperative systems – as *emergent* properties rather than as being imposed by some higher-order executive level distributing instructions to lower levels. Such an alternative viewpoint is not entertained by Arbib in the target article, so his theoretical approach has to be seen as an attempt to explain the actions of animals (and, ultimately, man) in computational terms utilizing the concept of perceptual and motor schemas. Whereas such an approach leads Arbib to raise the question "Are schemas real?" those (like Thelen) favoring a dynamical approach would want to pose the alternative, and more searching, question, "Are schemas necessary?" A rigid application of Occam's razor would suggest that they are not (Kugler 1986). Why invoke mental representations if more powerful explanations can be found for the same phenomena? Why introduce special mechanisms (e.g., motor schemas) when explanations of the same observations might be derived from first principles (Kelso & Scholz 1985)?

If schemas (motor) are real, as Arbib hints, where do they come from? Apart from distinguishing between episodic and skill learning, Arbib has little to say about learning or development – particularly as they relate to motor actions. It is hence interesting in this respect to take note of Prechtl's (1986) point that the neonatal movement repertoire does not appear *de novo* at birth but is preceded by a prenatal developmental course. Prechtl (1986), with the aid of real-time ultrasound, was able to show that fetal movements that later comprise the motor repertory of the full-term neonate are fully formed by the fifteenth to sixteenth week of gestation. Is it also necessary to invoke schemas for their motor control? If so, were they present prior to the movements or added at a later date?

This is not the place to pursue this debate further (the interested reader is referred to Meijer and Roth, 1987, for an explication); it is worth pointing out, however, that exponents of the "alternative approach" here referred to would be less concerned than Arbib with whether their interpretations might survive the tests of the psychological laboratory and more concerned with whether they survived the tests of everyday functionally oriented behavior.

Many kinds of criticism might be, and have been, leveled at the notion of schemas, but the particular problems in Arbib's target article are its vagueness and arbitrariness. A good example of the former is provided in the following sentence (sect. 2, para. 3): "We hypothesize that these plans are made up of *motor schemas,* which are *akin* [emphasis added] to control systems but can be combined to form coordinated control programs to control the phasing of various patterns of co-activation."

The motor schema turns out to be very difficult to grasp indeed, particularly when, as in Arbib's article, it is used in contexts as different as actions, behavior, and movements. Recourse to an earlier article of Arbib et al. (1985) on coordinated control programs for movements of the hand helps a little in confirming that the motor schema is in fact conceived of as "a unit of *motor* control," but the reader is still left to speculate on the precise nature of that control: "A motor schema is a control system, continually monitoring feedback from the system it controls to determine the appropriate *pattern of action* to achieve its goals" (emphasis added; Arbib et al. 1985, p. 111).

The same article illustrates how arbitrary schemas can be. A task analysis for "grasping a mug" (for which there is, apparently, also a schema) gives rise to the identification of a number of subactions: "From a resting position . . . the hand preshapes while the arm reaches . . . the grasp begins, and . . . the mug has been actually grasped" (Arbib et al. 1985, p. 112).

To account for the operationalization of this task, a *reach* schema and a *preshape* schema in addition to a *grasp* schema are postulated. Presumably, if the task analysis were to be further refined, additional schemas could be quickly invoked! [See also Berkinblit et al.: "Adaptability of Innate Motor Patterns and Motor Control" *BBS* 9(4) 1986.]

The problem with invoking ontologies, such as schemas and programs, is that unless one wishes to resort to serial chaining interpretations, another ontology has to be proposed (in Arbib's case coordinated control programs) to account for interactional parameters (in Arbib's case the time-varying interactions of a number of motor schemas) between sets of schemas (programs).

Bernstein (1935; see Whiting 1984) presents a paradigmatic example of the problems to which this can give rise. In trying to account for the regularity of rhythmical movement he insists that the observed homogeneity "originates in the operations of the central nervous system": ". . . there exist in the central nervous system exact formulae of movement (*Bewegingsformeln*) or their engrams, and that these formulae or engrams contain in some form of brain trace the whole process of the movement in its entire course in time" (Whiting 1984, pp. 97–98).

Bernstein's "proof" of the existence of such engrams was "the very fact of the existence of habits of movements and of automatized movements" (Whiting 1984, p. 98). His corollary was that there must be some guarantee of the order and the rhythm of the realization of such schemas. Bernstein's ontological solution was the *echphorator:* ". . . the echphorator lies outside the engrams themselves and directs their order by a hierarchic principle of *Überordnung*" (Whiting 1984, p. 98).

Thus, as Reed (1984) points out, a potential, the ability to act in a determinate way, is turned into an actuality (a latent engram) and the actuality is claimed to cause all instantiations of the potential. This kind of muddled thinking, he maintains, together with the assumption (to which all existing theories of indirect action seem to lead) that the brain causes the body to move and that actions are the environmental consequences of such movements, needs to be abandoned.

Author's Response

Of schemas, neural nets, and *Rana computatrix*

Michael A. Arbib
Departments of Computer Science & Neurobiology, University of Southern California, Los Angeles, Calif. 90089

The target article strove to combine the study of action and perception, the analysis of behavior at the molar and the neural levels, and the approaches of artificial intelligence, brain theory, and cognitive science. As a result, the commentators come from diverse fields. I have made a rough inventory of the commentators' perspectives in Table 1. Perspectives that are lacking include those of Piagetian psychology, linguistics, and the two "extremes" of analysis: the social level of knowledge and neuronal analysis in terms of detailed membrane or molecular mechanisms.

The structure of this response is as follows: The general theme is that the schema level of analysis is of interest even to those cognitive scientists who are not concerned with neural implementation. However, for the *neuroscientist*, schemas have utility in providing a functional description to constrain, and be constrained by, analysis into neural networks, as I indicated in the target article. Moreover, the study of *Rana computatrix* is not only a tool for the study of the neuroethology of frog and toad, but offers organizational principles useful in studying the vertebrate brain in general. In the sections that follow, we will note the agreement of most commentators as to the need for intermediate levels and then we will discuss the various claims as to what those levels might be. We will note that the "language of schemas" need not in itself be exhaustive as we ponder the explanatory power of adaptive networks and nonlinear systems. However, the need for schema-like units is not seen by everyone, so we must turn to a critique of Gibsonianism and emergent properties. The concept of schema is not a new one, and the diverse models of schema offered here are but stages in developing a schema theory. Some of the issues are discussed in the sections on probing the schema for schemas and on instantiation. As we move from the level of schemas as abstract entities of interest to all cognitive scientists to their possible implementation in the brain, we first look at neurological data at the level of brain regions, analyzing brain lesion data as constraints before turning to the more detailed level of neural networks. When it comes to neural networks, we first note the general implications of our study of *Rana computatrix* for computer experimentation and modeling style before turning to future directions in the modeling of frog and toad per se. In the final sections, we turn to the broader implications of schema theory by addressing explicit questions about cognitive psychology and about schemas and persons.

Table 1. *Inventory of commentators' perspectives*

Experimental research on frog brains: **Ewert, Matsumoto, Székely**
Neuropsychology: **Marshall, van Leeuwen**
Neural modeling: **Baird, Johannesma, Tang, von Seelen**
Schema theory: **Barnden, Mackworth, Tsotsos** (the last two in computer vision)
Philosophy: **Gunderson, Lloyd**
Cognitive psychology: **Langley**
Experimental psychology: **van Leeuwen, Wagemans**
Kinesiology: **Whiting**

Intermediate levels. The commentators agree on the need for many levels of analysis in cognitive science, but not on what those levels should be. It is interesting that no commentator extends the levels beyond the span from mind to neuron, neither going "up" to the social nor "down" to the neurochemical. Newell (1980; 1982) has argued for a symbolic level between the physical level and the knowledge level; Pylyshyn (1980; 1984) distinguishes between the level of basic mechanisms, a syntactic or computational level, and a semantic or representational level. **Wagemans** suggests that whereas a schema might be considered as occurring at the highest level in the mainstream of cognitive science I see it as an intermediate level: The content of the schema is determined by its formal computations and as such is not semantic in the strict sense. Perhaps we can reconcile these levels by distinguishing a schema as specified purely in terms of its functional input/output behavior (Pylyshyn's representational level?) from a schema whose function is specified in terms of internal information processing (Pylyshyn's computational level?). Again, if we ignore the numerical details essential to the function of perceptual and motor schemas and instead simply stress the names of schemas and their abstract relationships, we seem to be close to Newell's symbol level:

knowledge level
symbol level = schemas as abstract symbols
functional level = task level = schemas as input/output descriptions
computational level = add internal programs to schemas
connectionist level = implement programs as connectionist nets
neural level = implement programs as plausible neural networks

Tsotsos argues that, just as there are many different branches of mathematics suitable and appropriate for different problem domains, so should we expect a family of varied representations for levels of analysis of the brain. The target article certainly presents many different representations. We might also note that there are alternatives at the implementational level, such as nonlinear systems and associative nets. Thus **Mackworth**'s observation that I use four schema formalisms in my paper may not be the negative point it first appears to be. It is an open question whether the same formalisms are appropriate for perceptual schemas (e.g., Appendix A) as for motor schemas (e.g., Appendix B); whether activation levels are needed in the study of motor schemas (as argued by Dakin & Arbib 1986); and whether different schema formalisms are appropriate en route to, say, a LISP rather than a neural net implementation. The open research issue is to discover which differences are essential to the differences in task domain and which are historical accidents to be removed by the unification of different minitheories.

Gunderson feels that my claim that "schemas must be large enough to allow an analysis of behavior at or near the psychological level" suggests a level of behavioral analysis free of psychological ascriptions and one that inevitably uses them. In particular, he asks, what is the difference between a psychological level and a "near"-psychological level? What do the "large"-enough schemas grab onto? My use of the term "near-psychological" was not intended to suggest a level "free" of psychology. Rather, it was meant to suggest that the vocabulary we receive from psychology may need modification as the details of the implementation are worked out. This was the burden of my espousal of a "two-way reduction" whereby the reality of different levels may be recognized and reconciled to create a richer whole, but wherein the realities may also change. One example is the way in which the analysis of depth perception shifts from an unanalyzed psychological concept to a set of competencies that are "near" the psychological level – perhaps they can be said to redefine the psychological level – but that are precise enough to be explained in terms of finer-grained mechanisms.

It is in the same spirit that I would respond to the comments of **Mackworth** and **van Leeuwen** concerning Marr's notion of the "computational level" (which is better called the "functional level" or the "task level"). I did not reject the level, and in fact explicitly used the "functional level" or "task level" as my point of departure for the analysis of behavior long before Marr's discussion of it (cf. the stress on *action-oriented perception* in Arbib 1972). What I do reject is the claim that analysis at this level provides a description resistant to modification when we take "implementational details" into account. This is what I meant by saying that Marr's notion of the computational level as an *independent* level of analysis is mistaken. As **Wagemans** points out, Marr admits that a given computational task (e.g., of recovering depth from 2D images) can be done by several algorithms. This is true, but the point is that the *initial specification* of the task at this level may not survive the attempted implementation. Thus Marr placed great weight on the computational task of computing a $2\frac{1}{2}$D sketch prior to image interpretation, but few in the machine vision community now accept this as more than, at best, a partial "computational" specification. And to say that Marr worked with mathematical rigor and computational detail (which was certainly not the unique and *ab initio* creation that some believe) is not relevant to the evaluation of his doctrine of the computational level.

In continuing what he stresses are only minor criticisms, **Wagemans** asserts that I am wrong in characterizing modules as structural units, citing the authority of Marr (1982) and Fodor (1985). But neurophysiologists had already made use of the term in the 1970s (Szenthágothai & Arbib 1975; Mountcastle 1978) to denote a structural unit in the brain, and even earlier Kilmer et al. (1969) had used the term in modeling the reticular formation as a chain of interacting modules, each corresponding to a "poker chip" of neural tissue. Actually, I think that the word "module" is robust enough to be used with either adjective, "structural" or "functional," but I have argued elsewhere that there are major incoherencies in Fodor's account of what constitutes a (functional)module (Arbib, 1987a).

Adaptive networks and nonlinear systems. It is suggested by **Baird** that the discovery of the implementational details of a biological system may contribute to the way a task is conceived and described at the functional level. A network with a learning algorithm, such as Holland's (1975) classifier system, may be capable of generating many specific algorithms to solve large classes of problems and so may encompass large classes of symbolic algorithmic descriptions even when we do not know what specific symbolic description is appropriate. We may note too the recent success in training hidden units in "unstructured" networks to enable the nets to exhibit previously unspecified functions in adaptive reward-penalty (A_{R-P}) (Barto 1985) or back-propagation networks (Rumelhart et al. 1986). Still, the brain theorist has to go beyond such studies in trying to understand the structural particularities of different brain regions and their variation with behavior from species to species.

Baird agrees with my advocacy of the strategic importance of the incremental neural modeling of a specific well-chosen system, citing the work by Walter Freeman's group on pattern recognition in the rabbit olfactory bulb [see Skarda & Freeman: "Brains Make Chaos in Order to Make Sense of the World," *BBS* 10(2), 1987]. They use an 8×8 array of electrodes to monitor spatial patterns in the EEG during the conditioning of a simple licking response to water reward. The data lead them to analyze pattern recognition networks that bifurcate from a single background state into a state space with multiple periodic and chaotic attractors (Baird 1986a; 1986b). In fact, there is now an increasing interest in the use of chaos to model apparently random behavior in neural systems (cf. Freeman & Skarda 1985; Kelso & Scholz 1985; Mpitsos & Cohan 1986). Applications of nonlinear analysis by my own colleagues include the study of competition and cooperation in neural nets by Amari and Arbib (1977) and the parameter analysis of the tectal column by Cervantes-Perez (1985, Chapter 4).

Gibsonianism and emergent properties. My work is seen by **Wagemans** as connecting divergent scientific strategies (theory building vs. experimental testing), between metatheoretical approaches (ecological vs. constructivist) and between research domains (perception and action). He notes that Gibson (e.g., 1979) would have called schemas, or any other intervening processes or structures, superfluous. Wagemans sees my definition of a perceptual schema as "a process whereby the system determines whether a given domain of interaction is present in the environment" as compatible with the ecological approach, but much more explicit as to how the tuning of the system to environmental affordances actually occurs. I have indeed found Gibson very helpful in specifying facets of perception that are worthy of study, but I then move beyond him in insisting that we cannot just pick up affordances without some mechanism (e.g., optic flow; Prager & Arbib 1982).

Whiting adopts a strongly Gibsonian position that seems to leave little room for agreement. He notes that Thelen's (1986) double perspective on movement in the time domains of both actual movement and development leads her to interpret the observable in each domain as an *emergent* property rather than as being imposed by some higher-order executive level distributing instructions to lower levels. Since I do not discuss this viewpoint, Whiting asks, "Are schemas necessary?" Why introduce special mechanisms (e.g., motor schemas) when explanations of the same observations might be derived from first principles (Kelso & Scholz 1985; Kugler 1986)? But this may be a confusion. A theory of schemas must include a theorv of schema formation, and the work of people like Kelso may let us understand how a schema could emerge from the cooperative action of a network. What has emerged is still worthy of cataloguing, however, and its interaction with other motor schemas (as I shall continue to call them) still needs analysis.

Whiting cites Arbib et al. (1985) to note with disapproval how arbitrary schemas can be. Yet most cognitive scientists would agree that "something" is required to tell a mug from a screwdriver, and that, arbitrary though the distinction may be, picking up the mug and using the screwdriver require different skills. Many kinesiologists agree that complex skills benefit from an analysis into interacting components. I call these components motor schemas, as Schmidt (1975) does, for example, and I am developing the language of coordinated control programs to make explicit how they can be combined. Whiting seems to feel that to call them "emergent properties" renders explanation superogatory. We are back to Gibson again – but direct action instead of direct perception. He applies to my coordinated control programs Reed's (1984) critique of Bernstein [see also Berkinblit et al.: "Adaptability of Innate Motor Patterns and Motor Control Mechanisms" *BBS* 9(4) 1986] for turning an ability to act in a determinate way into an actuality (a latent engram), with that actuality claimed to cause all instantiations of the potential. On the contrary, my discussion of Piaget in Arbib (1985), while accepting the validity of schemas, stresses that the problem of telling which schema a given behavior falls under is nontrivial. Reed maintains regarding schemas as actualities is "muddled thinking" and needs to be abandoned, along with the assumption (to which, he remarks, all existing theories of indirect action seem to lead) that the brain causes the body to move and that actions are the environmental consequence of such movements!

Clearly, anyone committed to understanding the neural basis of movement cannot deny that the brain causes the body to move and that actions are the environmental consequences of such movements. However, I do take the following points: There need not be a coordinated control program separate from the interaction of neural circuitry embodying the constituent motor schemas; that is, the coordinated control program is a description, in this case, of transfers of activity in neural nets, not a pattern in a separate neural net that gates those transfers. Nonetheless, the data of Evarts and Tanji (1976) do seem to point to the neural reality of such gating in some circumstances, while our ability to acquire the rudiments of a skill through verbal instruction seems to point to the ability to assemble schemas. This is not to deny that once a high-level specification has "homed in" on an individual schema, the more detailed properties of that schema may thereafter be better understood in terms of cooperative properties of neural nets than in terms of symbolic descriptions. Here we may recall the comments of **Baird**, and profit from the insights of Kelso. I must reiterate the point made in my commentary on Anderson's accompanying target article [this issue]: In computer science, the translation from a program in a high-level language to its machine language implementation preserves semantics, whereas in cognitive (neuro)science, the high-level description may only approximate the behavior of the fine-grained implementation.

Tang seems to share some of **Whiting**'s concerns when she argues that it is important to find more methods for representing schemas. Whereas "frames" represent well-learned knowledge in AI and are suitable for the analytic mode of thinking that breaks the world up into parts and then combines them, she speaks of another mode, *imagistic thinking,* which processes the world as a whole. Noting that children abstract the concept of cat from one or a few cats and can later recognize any cat immediately, she asks, "What schema representation will facilitate such analogies?" There are two immediate condidates. One is schemas implemented as associative networks (Kohonen 1984). The other is schemas of the "analytic" kind, but supplemented by the sort of inferential mechanisms provided by Indurkhya (1986) in his theory of metaphor (and see the schema-based treatment of language as metaphor in Arbib and Hesse 1986). However, Tang does not follow Whiting and Reed in their rejection of brain-based studies of movement. She suggests that it would be useful to have a hierarchy of subnetworks in a schema, wherein schemas located at the same level would interact by competition and cooperation, whereas schemas at different levels would act by maintaining the proper sequence. For example, she sees motor control from the motor cortex to the cerebellum, brain stem, spinal cord, and muscle-servo as clearly hierarchical.

Probing the schema for schemas. It is argued by **Tsotsos** that there is no single required level of representation between behavior and brain, and that schemas do not yet provide the desired interlingua. The designer of a representation must determine how appropriate it is for a given problem domain and must provide a syntax and formal semantics so that others may adopt it as well. **Mackworth** notes that since the modern introduction of the concept of schema by Head and Holmes (1911) the schema for a schema has been "extraordinarily plastic." He notes that, although I provide some functional criteria for schemas and four structural definitions, the definitions are far from coextensive and thus schema theory in its current state cannot serve as the common language of cognitive science. I have already stated, in agreement with Tsotsos, that cognitive science may be better served by a variety of schema formalisms; but I have also confessed that present formalisms (both mine and those of others) are provisional. We may thus expect future work to benefit from Mackworth's observation that functional analysis at the task level can lead to the identification of adequacy criteria that *any* representation scheme must satisfy – he has identified eleven descriptive and five procedural adequacy criteria that can serve as functional constraints for schema theory (Mackworth 1987).

As **Barnden** and **Mackworth** note, the body of a schema is specified in Appendices A and B by means of a fairly conventional analogue to current programming languages, with numerical parameters, sequential control with program counters, and minor modifications for parallelism, including the FORALL mechanism and coroutines (for related constructs cf. Tsotsos 1985; 1987). To put this in perspective we must distinguish between a language for specifying schemas at the schema level, however imperfect (we are still learning languages for high-level, possibly distributed, function), and the neural implementation. The formalisms of Appendices A and B were designed for ease of AI applications on current computers. The Lara example comes closer to my view of schemas-in-the-brain and their implementation on future neural computers. The key concepts are port connections, instantiation, deinstantiation, and activity levels mediating competition and cooperation. Barnden suggests allowing the behavior of a schema to be specified by means of a mathematical description of an input–output function or relationship. This is what is accomplished by the Lyons and Arbib (in press) work using port automata (Steenstrup et al. 1982) for the semantics of schema instantiations. As with automata, this may be specified by an explicit input–output history or implicitly by state-dependence. Barnden adds that it may be relevant to know for each port of a schema the style of encoding and the required degree of precision for the data it carries. This could de done by extending Lyons's (Lyons & Arbib, in press) association of a data type with each port.

Tsotsos finds a host of questions raised by my treatment of schemas. These I shall now attempt briefly to address:

1. *How are schemas mapped onto an implementation for testing purposes? Is there a translator?* There are translators for the schema systems of Appendices A and B that map schemas into programs for current computers. However, although neural nets do have a syntax and semantics, there is no systematic way to map schemas to neural nets when we are studying the brains of actual animals. In fact, as I stated above, the descriptions at the schema level and the neural level can be expected to coevolve, one constraining the other, when we are doing neuroethology.

2. *How can schemas proposed for the same behavior be compared without simulation?* Automata theory certainly offers techniques for proving that two automata do or do not exhibit the same behavior, and these methods can be extended to port automata and thus to schemas. At the neural level, Cervantes-Perez (1985) shows how to prove that a given range of parameter settings will yield a certain qualitative class of behaviors. **Mackworth** notes that there are intriguing new results from theoretical computer science to be exploited. Complexity analysis shows that some tasks are inherently serial, and Mackworth (in press) has shown that establishing a form of local coherence (arc consistency) in a set of declarative schema instances is such a problem. Nonetheless, a good simulation test-bed will remain indispensable for the investigation of schemas and neural nets.

3. *What class of behaviors does a given schema represent? How can one tell whether that class of behaviors is the desired one, and contains no unwanted behaviors? How can one tell whether a given schema formalism is sufficient for all the behaviors of interest and only those?* My own suspicion is that any interesting class of schemas will have sufficient expressive power to compute all computable functions. The interesting questions then will be to identify schema networks that can match data on speed (cf. **Mackworth**'s complexity theory), errors, and lesions. Chomsky has sought to characterize a class of grammars that can represent all and only potential human languages, but the current formalism can certainly represent unnatural languages as well. [See also Chomsky: "Rules and Representations," *BBS* 3(1) 1980.] Amari and Arbib (1977) and an der Heiden (1980) are among the many to use a leaky integrator model of the neuron, but Hopfield and Tank (1986) have shown that one can map the "travelling salesman problem" onto an unnatural neural net. Thus I do not think we will see schema formalisms that represent only those behaviors of interest to, say, the linguist or the neuroethologist.

4. *Is a schema a procedural or a declarative entity? Does it matter? Are the arrows between units within a schema control lines or data lines, or do they represent some other relationship?* At present, I think of schema instantiations as actively interacting entities. In the coordinated control program for grasping I have specific lines for transfer of control and data. Appendices A and B make clear the need for passage of activation levels and (de)instantiation commands. **Tsotsos** notes that Hayes (1974; 1979), Woods (1975), and Brachman (1983; 1985) directed attention to the aim of providing a representational formalism within a logical semantics, but to my taste their style of semantics marches us off in the wrong direction, away from what might be called a "procedural cooperative computation" approach to schemas that can offer a bridge to connectionism or to neural nets.

5. *What are the primitives of schemas? Is it possible to define a set of primitives so that all behaviors of interest may be represented using combinations of these primitives?* In a particular application domain (such as in Appendix B), we may work with a fixed set of basic schemas. In a general setting, I see new schemas being formed as assemblages of old schemas; but once formed a schema may be tuned (by something akin to an adaptive network mechanism), much as a skill is honed into a unified whole from constituent pieces. It is the tunability of schema-assemblages that allows them to start as composite but to emerge as primitive. It is also the reason why I stress that a model expressed in a schema-level formalism may only approximate the behavior of a model expressed in a neural net formalism. (Perhaps this is what **Whiting** had in mind. Knowing that a well-tuned skill may not be decomposable into pieces, he rejects the role of motor schemas in acquisition, and in more structured skills.)

6. *In motor tasks, temporal issues are of particular importance. How is time represented in the schemas?* See Fuster (1985) and Tsotsos (1987).

7. *How are motion qualities such as "grasp gently" or "grasp aggressively" represented? What are the subtleties of vision that current schemas cannot capture?* I do not know.

Langley finds the representation of motor schemas (in

Appendix B?) to be implausible, being too influenced by the use of numeric representations and algorithmic control structures in robotics. Well, it was a robot programming language! Langley seems to make the mistake of many cognitive scientists who have not sufficiently pondered low-level vision or motor skills; they ignore the parametric nature of much human behavior. As we develop schemas tuned to human motor control, as in the work of opposition spaces of Iberall et al. (1986), even if the schemas as wholes act as symbols (see the above discussion of Newell's [1980] symbol level), the internal structure of a schema preshaping a hand, say, must have properties more akin to numeric representations. Langley cannot believe that this schema language will give the flexibility and adaptivity required to model human motor behavior, nor does he see how it would support models of motor skill improvement and acquisition. In fact, as indicated in my response (5) to **Tsotsos,** I see the explicit language of schema combination as pointing to the way in which we can acquire the rudiments of a skill by verbal instruction; but it is only with something like an adaptive network implementation (cf. **Baird**) that the honing of a skill can be explained. The PASCAL-like program is one way of specifying the network top-down before we refine it through network analysis.

Instantiations. The notion of multiple instantiations is seen by **Barnden** as central, presenting deep problems when the task of implementing the schema level in neural circuitry is considered. However, my paper does not involve *detailed* attention to multiple instantiations, and there is no consideration of how multiple instantiations would be neurally realized. Since I agree with Barnden as to the importance of these issues, I use this section to emphasize this direction for future research.

Barnden notes that one connectionist solution to the problem of cross-talk between similar pieces of information is to recruit neural assemblies to represent the particular instances of a situation class, as well as to have a neural assembly standing for the situation class itself (e.g., Hinton 1981). But how is short-term recruitment managed? Does it rely on synaptic-weight change? How can it be made fast enough? How economically can inferential and other mechanisms respond to recruited – and so in some sense unpredictable – assemblies? How are recruited assemblies demobilized? This is especially difficult when an assembly shares many neurons with other assemblies. Von der Malsburg (1981; von der Malsburg & Schneider 1986) offers a solution to some of these problems with his dynamic link architecture. He uses synaptic change on a long time scale for long-term memory; fast modulation of synaptic weights about the long-term values then provides the dynamic relationships within and across assemblies. This architecture appears to be promising for neural instantiation of schemas.

Can we produce a theory that has different instantiations of a schema S as superimposed states of a single neural net N? Or are the different instantiations different copies of N dynamically recruited, or perhaps permanently existing but only intermittently active? **Barnden** suggests that we contrast an N that is specific to a given S with an N that can be programmed to accommodate a wide variety of given S's (as in the universality of cellular automata). Or might the instantiations reside in networks that copy only parts of N? Here we may compare Weymouth's (1986) parts (Appendix A) versus Lyons's (1986) wholes (Appendix B). And how does all this relate to the number of instantiations of a given schema that can be simultaneously active? Perhaps only instantiations are identifiable with neural circuits, whereas a schema might be just a propensity to activate neural circuits to act as its instantiations. In this connection, it is tempting to recall Heinzel and Selverston's (1985) circuits, which serve different motor schemas (CPGs) under proctolin control.

In an earlier issue of *BBS*, Arbib and Caplan (1979) discussed the DIPM model (Arbib 1970) in the context of the effects of brain lesions. The idea was that a schema could set up instantiations in different parts of a cortical network, but with time it would have a preferred subnetwork in which appropriate parameters were stored. Lesions in this specific region, then, would not irreparably destroy the schema, but would greatly reduce the skill with which it could be evoked. Similarly, in his research on connectionist models for complex short-term inferential information-processing, **Barnden** (1985; 1986) maps abstract connectivity among pieces of information to an intermediate level in which there is a notion of relative positions of data items in certain representational media. These relative positions relate to hardware connectivity in a complex way.

Brain lesions as constraints. Analysis at the task level, as **Mackworth** observes, allows us to determine the natural constraints that *may* be exploited by a system and the structure of the equivalence classes of solutions to the perceptual or motor task. We can determine whether a solution is overconstrained or underconstrained by the various potential information sources. If it is underconstrained, then additional constraints or preferences can be identified that may be imposed to select out a particular solution (Mackworth 1983). I agree, but my point is that these constraints may be either at the task level or at the implementational level – as in the constraints revealed by lesion studies.

Marshall provides a fascination catalogue of data on humans with brain lesions that dissociate different aspects of visual behavior but he confesses that he does not know whether *Rana computatrix* will play a heuristic role in elucidating disorders. Instead, he expresses the wish that "more members of the simulation community would take a look at the kinds of data provided by disorders of human cognition." Of use here is the book edited by Arbib, Caplan, and Marshall (1982) in which I have three chapters that build upon the *BBS* target article of Arbib and Caplan (1979) to attempt to apply schema theory to the neurology of human cognition. Marshall confesses that my observation that different animals "may make different use of visual cues" causes him to worry whether *Rana computatrix* does it "our way." But I think this is to miss two points. First, my article offers schema theory as a general methodology for cognitive analysis. For neuropsychology, we seek schemas that can be localized in the brain by lesion studies, and this has nothing to do with *Rana computatrix*. The second point, however, is that we *can* learn about the human brain from *Rana computatrix*, but not by simply plugging frog schemas into the human brain. *Of course* the visual systems of frogs and humans are different – this is already dramatically evident at the level of the retina. My claim is not that a model of the frog is also a model of the human; it is rather that the study of *Rana computatrix* enriches the vocabulary of *organizational principles* and *model exemplars* with which we may approach the specific problems of human cognition. For example, Arbib (1982c) contains a section, "From Prey Selection to Object Naming," which uses a model of the neural circuitry underlying prey-selection in the frog to offer novel hypotheses about Luria's (1973) observations on how certain lesions affect a patient's ability to name objects. It is at this level that I would show the relationship Marshall desires between my kind of modeling and the problems his patients present. The appropriate models here are at the level of schemas as constrained by neurological lesion data, rather than at the level of detailed neural circuitry.

Van Leeuwen argues that constraints on representations are at least as fundamental as those on processes, and with this I certainly agree. But he then states that this means that the constraints themselves should be formulated independently of any semantic domain, and that schemas are domain-dependent and therefore cannot be expected to serve as a basis for understanding the neurological constraints on computational theories. I do not agree with this logic. The neurological constraints enter when we have data on the effects of brain lesions (though these are complicated by the brain's ability to adapt), or data from PET-scans or evoked potentials. These constraints cannot be brought to bear unless the schema model is refined to a level where specific hypotheses can be made as to the localization of subschemas in the brain. As van Leeuwen indicates, however, they can also constrain the schemas by requiring that those which survive a lesion be able to account for the representations then exhibited by the patient. This is what **Marshall** was doing for us when he demonstrated that different brain lesions could fragment representations that might otherwise have appeared unitary. We are then challenged to work with a mosaic representation consistent with these data (cf. the accommodation vs. disparity of the cue interaction model), rather than with the unitary representation that tempted us in the first place. This is precisely why I argue that the functional level must be accountable to the implementational level – our a priori choice of representation may not survive the provision of constraints from lesion data or other "implementational levels."

Computer experimentation and modeling style. The commentary of **Langley** expresses the faith that it is possible to offer functional analyses that have value despite all lack of neural grounding. How might this be? Perhaps because different schemas may be functionally equivalent with respect to such criteria as predicting a frog's behavior or teaching a child a new skill (cf. Anderson's target article in this issue). For many cognitive scientists, such a grain of explanation may suffice. For the neuroscientist, however, schemas are to provide hypotheses as to the function of specific neural circuits; these are falsifiable; new hypotheses (and the record of mistaken hypotheses) then lead on to new understanding. **Lloyd** states that a complex neural model should be *complete* in two senses: The behavioral capacities should be given neurally plausible implementations, and the pathways modeled should extend without break from sensory input to motor input. Since the depth and detour studies are still mainly top-down, they may be undermined by later data on localized brain functions.

Since there are many parameters to be tuned in the tectal column model, **Lloyd** wonders what guides the choice of adjustments. Simplifications may have ramifying effects: Increments have more than incremental consequences. This is true, and I regard future attempts to combine tectal column modeling with depth and detour modeling as posing a major challenge to both models. I do not have an algorithm to guide this merger. Initially ad hoc methods may be required to get the new model to meet a barrage of constraints provided by experimental data. Lloyd suggests that there is a trade-off between accuracy and completeness and argues that, if one must choose, completeness is the more useful goal. I sympathize with the move to completeness, stressing visuomotor coordination rather than vision per se, but I must note that many experimentalists focus successfully on specific regions in the brain, rather than following through to each periphery. Thus, much modeling designed to make contact with neuroscientific data must sacrifice completeness for accuracy.

It is in the spirit of his call for completeness that **Lloyd** finds *Rana computatrix* of cognitive interest, irrespective of its accuracy as a model of living anurans. Dennett (1978) proposed that one approach the complexity of humans by looking at simpler systems – both living systems, and artefactual systems born of engineering imagination. Braitenberg's (1984) "vehicles" exemplify the latter approach, a synthetic psychology, and Lloyd suggests that *Rana computatrix* be located in the same phylum of *Gedanken* beasts, "the most advanced of its kind yet to evolve." Dewdney (1987) has also made the point that *Rana computatrix* should be seen as a sophisticated "vehicle" in Braitenberg's sense.[1]

Noting that I view the animal's internal model of the world as an assemblage of instantiated schemas that correspond to "domains of interaction," **Lloyd** asks whether the instantiated schemas in the brain of *Rana computatrix* represent flies, small distal objects, or small patches of retinal stimulation? Similarly, under what conditions will we say that *Rana computatrix*'s model of the world is false? These are central issues for cognitive science, and complete models like mine can help their resolution (cf. Lloyd 1987). Such questions are addressed in my discussion of pattern recognition in my commentary on **Ewert**'s accompanying target article. I suspect that a model of the world is not so much true or false as it is more or less useful – the frog's model is useful if it gets enough to eat, and does not expend too much energy snapping at false flies. [See also Dennett's accompanying commentary on Ewert, this issue.]

I think such concerns prompt **Johannesma**'s commentary, but he seems to reject completeness, offering instead what appears to be a methodology for considering a sensory system in isolation. Johannesma asks how *Rana computatrix* may gain access to the biotope of *Rana naturalis*. I thought I had done this by offering simulations (such as the Lara et al. 1984 model) that went all the way from visual input to behavioral output. Thus, though I find Johannesma's ideas of interest, I am not sure they advance the notion of task-analysis or action-oriented perception, which seems central to his "sharing of the biotope." He analyzes neural representation of sensory activity as a map P from sensory space into neural space, and the sensory interpretation as the inverse map Π of neural space into sensory space. We must immediately note that there may not be a unique inverse. This need not be a problem for Johannesma, however, who stresses probabilistic interpretations, for he later asserts that if sensory stimulus x induces the activity pattern z, $z = Px$, and y is the sensory interpretation of z, $y = \Pi z$, then the sensory interpretation of x is $y = \Pi Px$, where ΠP is "a deterministic or stochastic mapping of sensory space onto itself." However, we should note here the well-known fact that Π is so underdetermined by P that extra constraints must be imposed if ΠP is to be a useful interpretation.

TRANSFORMATION

$$\text{sensory space} \xrightarrow{P} \text{neural space} \xrightarrow{Q} \text{motor space}$$

REPRESENTATION

$$\text{sensory space} \xrightarrow{\Pi} \text{neural space} \xrightarrow{Q} \text{motor space}$$

Johannesma states that transformations P and Q can be divided into subtransformations or schemas, but I do not think this quite captures the full richness of a system that includes feedback, relaxation, or other cooperative and competitive processes. Johannesma asserts that the formal description of the nervous system should be in terms of the forward mappings P and Q, whereas the understanding of function should be based on the representation given by Π and Q. Again, I do not see why P and Q (and thus the simulation of animal behavior) are less functional than Π and Q.

In the cue interaction model, the activity of the accommodation field and the interaction field form an internal variable of *Rana computatrix*. **Johannesma** asserts that no appropriate metric has been or can be defined to compare internal variables of model and animal. Even neural variables of different animals cannot be compared directly. But neurophysiologists certainly compare classes of receptive fields across animals, as in the work of **Ewert** cited in my target article. Application of a sensory interpretation of the internal variables by an inverse map of neural space into sensory space would result in a structure of the visual space of *Rana computatrix*, which then could be compared to the visual space of *Rana naturalis*. The map ΠP sets limits to the frog's perception. The map Π allows an external observer to compare perceptual abilities of *Rana computatrix* and *Rana naturalis*. Artificial and natural frogs would then be compared not with respect to their internal variables, but as companions sharing a common biotope. But as already stated, I think our action-oriented approach does create a common biotope. Correlation of internal variables then provides a structural/process constraint.

Johannesma does seem to be addressing a crucial question for the sensory neurophysiologist, however: How do we make sense of neural activity in sensory systems? His answer, in a symbol, is Π. On the basis of studies with his coworkers, he states that in sensory parts of the nervous system Π can be derived from the response of single neurons to sensory stimuli usually in terms of the spatio-spectral-temporal receptive field of the neuron. He further states that, since receptive fields are defined in sensory space, Π, being approximated as a superposition of the receptive fields of active neurons, will be defined as a distribution in sensory space. However, one must set against this the caution of **Baird** that a cell may exhibit very different receptive fields when in different behavior states. He suggests that single-cell "receptive field" or "trigger feature" investigations can be misleading or even incoherent when a network is cross-correlating its inputs. In his work on the olfactory system, he finds that apparent receptive fields may exist in such a net in a restricted experimental situation in which it always relaxes to the same attractor, but experimental conditions may change the attractor, and thus the relation of unit activity to input may be completely altered. Can Johannesma's methodology be adapted to the separation of these different maps Π_a, one for each attractor a?

Székely notes that even the experimentalist makes conceptual models when he describes and interprets the results of his study of so complex a structure as the nervous system. He thus warns the modeler that the tectal column of Székely and Lázár (1976) is but one of the many possibilities as to how the neuronal interconnections may be organized in the optic tectum. Similarly, **von Seelen** questions the column structure of neuronal networks, asserting that it has not been detected neuroanatomically and may be misleading functionally. In a layered structure, neurophysiological recordings suggest columns, but feedback within a layer or between layers negates the functional limitations of columns. Von Seelen prefers to view a neuronal layer as a basic system complemented – if needed – by discrete, geometrically arranged grids to which input is directed or where specific internal processing takes place (e.g., the "blobs" in visual cortex). The various types of network must be well analyzed and implemented as completely as possible. Investigation of the coordinate transformation of the retina onto the tectum may considerably simplify processing, especially with moving stimuli.

Székely urges the modeler to become a *computer experimentalist* to see what networks other than those sanctioned by current experiments can give better or worse results in terms of some arbitrary criteria. He can then point out to the biological experimentalist his faulty conclusions, overlooked data, or "emergent" properties that remain unnoticed in biological experiments. My discussion of the cue interaction model does in fact provide "attentional input" of this kind for the experimentalist. A model need not be a compendium of every shred of available data to be of conceptual value in helping us understand a class of neural circuits or in guiding new experiments. Székely finds the 8×8 array model of tectum "disappointingly trivial in showing the retino-pretectal-tectal interaction exactly in the same way as Ewert would unambiguously describe in three sentences." But Ewert's description is *not* unambiguous – it leaves a wide range of parameters unconstrained (recall **Mackworth**'s discussion of constraints). It takes intense computer experiments to find parameter ranges that fit the data, whereas, as **Ewert** notes in his commentary, other ranges do not fit. A given model network contains many different models depending on the choice of parameters. Cervantes-Perez (1985) showed that for specific parameter settings in our model changes in velocity could yield inversion of W/A (worm/antiworm) preference, a phenomenon not observed by Ewert and others. The interaction between theory and experiment thus allows one to constrain the operating range of certain parameters of the model. This variability of function within a given structure, even after the loci of excitation and inhibition have been determined, is one that few experimentalists appreciate, and is a vital one when we turn to neural modeling. For this reason, Ewert points out that the modeler can provide assumptions and approximations that are a source of hypotheses and predictions for the experimentalist.

Von Seelen illuminates the insights that may be gleaned from *Rana computatrix* for the general study of neural mechanisms of cognition when he asks: Why model simple toads that do simple things in simple experiments? His answer is that, given the conditions of the real world, the equations of physics are approximations that are frequently inadequate when boundary conditions occur. To describe brain functions, a certain level of abstraction has to be determined. He approves of my choice of basic schemas, which have to be constructible in neuronal architecture that can be measured with technical means, and then goes on to suggest that, as a help in finding and defining basic schemas, they should be interpreted as "basic situations," sensory or motor cues, or both, that the animal must perceive and handle. More complex behavior can be generated by combining basic situations. The spatiotemporal behavior of layered feedback structures can be used to encode and combine such representations. This approach assumes that brains are not universal computers but rather systems that have to solve concrete problems in a specific environment, that must be capable of evolution, and that can be constructed reliably under the constraints of ontogeny. The development of one new schema allows a large number of combinations with already existing ones. With all this I certainly agree.

Although advocating this approach, **von Seelen** notes that it raises problems, to which I now attempt brief answers:

1. *How can basic schemas or situations be defined and systematically detected?* We seek to model a complex behavior in terms of interactions among more simple behaviors represented by schemas. As I showed in Section 3.1, such a decomposition is successful to the extent that it passes the test of, for example, lesion studies. I do not know of any methodology guaranteed to yield the "right" decomposition, however.

2. *What strategy is used to couple the basic schemas when a target function has to be defined? Is there a general and workable data format for such systems?* As pointed out in the section on probing the schema for schemas, above, we now have several formal languages (e.g., Appendices A and B) for composing schemas into schema-assemblages and coordinated control programs, but these are subject to great improvement as a result of future research.

3. *Is self-organization possible?* Yes, although I have not addressed it in the target article. Such organization may involve adaptation either at the level of schemas and their interactions or at that of the neural or connectionist networks that implement the individual schemas (cf. the discussion of program synthesis and visuomotor coordination in the closing section of Arbib 1981).

Frogs and toads revisited. We have discussed the utility of schemas as providing a level of analysis that is useful to cognitive scientists even if they do not concern themselves with neural or connectionistic network implementations. We saw that *Rana computatrix* serves as a testbed for the understanding of organizational principles both at the level of schemas and at the level of (in many cases layered) networks. However, a number of commentaries focus on *Rana computatrix* not for its broader implications but to weigh its success as a model of the neural mechanisms underlying visuomotor coordination in the frog and toad per se. It is clear (as was already pointed out in the target article) that many new data have been accumulated of late, and that the time is now ripe for a vigorous round of new modeling.

Matsumoto notes that the original tectal column circuit was based on anatomical studies. He summarizes subsequent physiological data, which will surely contribute to further modeling. He is concerned that having dendritic endings of different cell types in the same glomerulus may be a poor simplification when discussing the time sequence of excitatory and inhibitory potentials. Whereas the original tectal column excluded large ganglionic neurons, Matsumoto et al. (1986) show that these neurons play an important physiological role in the tectum and exhibit a wide variety of morphology and physiology.

Matsumoto and Bando (1978; 1980) have observed interaction between myelinated and unmyelinated fiber systems in many tectal neurons. They identify the prey-selective efferent neuron as the pyramidal neuron and find that it receives a strong inhibitory input from class 3 retinal fibers. As in Ewert (1984), the current *Rana computatrix* involves thalamic inhibition, but no data are yet available as to whether the thalamic neurons directly inhibit pyramidal neurons or whether they do so through tectal interneurons.

If the optic nerve is electrically stimulated, large pear cells produce a monosynaptic and disynaptic EPSP (excitatory postsynaptic potential) followed by an IPSP (inhibitory postsynaptic potential). Matsumoto and Bando (1980) recorded the same response from large ganglionic neurons and found that the presynaptic fibers were myelinated. When large ganglionic cells were visually stimulated, some showed reverberatory responses (Matsumoto et al. 1986). If the stimulus suddenly stops in the receptive field, EPSP amplitudes are greatly enhanced, probably because of focal stimulation of the excitatory receptive field. This effect was greatest for a wormlike stimulus, which produces the smallest inhibitory effect on the neuron. We do not know, however, whether the reverberatory response seen in large ganglionic neurons is dependent on the positive feedback circuit in the tectal column.

Even in its present state, *Rana computatrix* calls for more detailed correlative morphological investigations, such as those cited in **Matsumoto**'s commentary. In his laboratory, **Székely** and colleagues are using serial electron microscopic investigations and immunohistochemistry to trace neuronal connections. Lara et al. (1982) have already conducted the kind of computer experiments he advocates on, for example, the attribution of inhibitory characteristics to different neuron types, but the simulation of the modulatory effects of various peptides has yet to be attempted.

One of the most exciting developments in the study of the frog is the increasing availability of experimental data on the motor systems, as exemplified in the *BBS* treatment by Berkinblit et al. (1986) on the spinal mechanisms of the wiping reflex. We also begin to learn something of the sensorimotor codes in specific bulbar/spinal motor pattern generators (MPGs) (cf. Figure 22 in **Ewert**'s target article). The discovery (Ewert et al., submitted; Schwippert & Ewert, submitted) of bulbar neurons displaying T5(4), T4, and cyclic bursting characteristics suggests interactions among highly integrative cells incorporated in a relatively compact "reticular" network. Ewert raises the question of whether different MPGs share general bulbar circuitry, depending on input from different "command releasing systems."

In addition to the pretectal structure, two areas in the anterior thalamus and one in the mesencephalic tegmentum receive optic fibers. Except for the latter center, all other areas receive a precise retinotopic projection of the visual field and are reciprocally interconnected in a remarkably organized fashion (Lázár 1984; Székely 1971). My target article does take some account of the highly organized reciprocal interconnections between the tectum and the nucleus isthmi. **Székely** would also like to see the incremental evolution of *Rana computatrix* help us understand the role of such highly organized interconnections in the control of visually guided behavior. Thus, *pace* **Barnden,** the theme of retinotopy is by no means exhausted.

Baird regrets that I did not discuss neural models of motor or motivational structures. Motivational effects are modeled by Cervantes-Perez et al. (1985; cf. Lieblich & Arbib 1982). He also laments the lack of schemas mapping cooperatively discriminated features to motor output, suggesting that there is mainly topographic mapping and local neighborhood interactions. In his system, preprocessing is followed by an associative network with global cross-correlating interconnections for object recognition. Might not frog models need nontopographic mapping to map object features to behaviorally, but not structurally, related patterns of motor activation? In fact, the model of Arbib and House (1987) allows widespread interactions in determining the path of the toad, even though those interactions take place in a retinotopic network. Boylls (1975; 1976) offers an interestingly different coordinate system for a network when he notes that the cerebellum is divided into microzones and suggests that each microzone is involved in setting parameters for a different motor schema.

With this, we leave the realm of neuroscience, and turn to the wider implications of our work, for cognitive psychology and philopsophy, respectively.

Cognitive psychology. Whereas the classic ethology of von Uexküll, Tinbergen, and Lorenz studied behavior in terms of single motor schemas, activated by specific perceptual schemas (innate releasing patterns), schema theory allows us to study the interpretation of the environment in terms of the *interactions* of perceptual schemas, yielding appropriate coordinated control programs of motor schemas. With the analysis of neuronal implementation, these schemas share the advantage of neuroethology as against classic ethology: We can come to understand how the evolution of the brain constrains the network implementation of schemas. Similarly, incremental modeling gives insight into brain mechanisms as we study how to change the model so it can accommodate the function of more and more schemas and their interaction. [See also Hoyle: "The Scope of Neuroethology" *BBS* 7(3) 1984.]

Newell and Simon's (Newell et al. 1958) original arguments for the computer simulation of human behavior focused on the notion of a *functional* description, and **Langley** sees me as reintroducing that concept in justification of schema theory. He also sees me as arguing for the linking of neural-level and functional-level explanations, but he himself doubts that this is the best path for cognitive science at this stage of its development. My abstract states: "Intermediate constructs are required as bridges between complex behaviors and realistic models of neural circuitry. For cognitive scientists in general, schemas are the appropriate functional units; brain theorists can work with neural layers as units intermediate between structures subserving schemas and small neural circuits." Perhaps this was not optimally phrased. My intent was to argue that schemas can provide a bridge from task analysis to neural analysis, but not to imply that the user of schemas is committed to the study of their neural implementation. I wanted to make two points: First, that schemas provide the appropriate functional language for *all* cognitive science; and second, that in *those* areas of cognitive science for which a linkage to neurophysiology is appropriate – vision and motor control are prime candidates – we may refine such a schema analysis with data on brain lesions (recall the above section on brain lesion data as constraints). Langley argues that we should not ignore the mechanisms involved in problem solving and natural language processing simply because we cannot yet explain them in neural terms, and of course I agree. This is the strategy taken in Arbib et al (1987). Meanwhile, note that the connectionists are showing the relevance of parallel processing for the study of a number of language problems for which the search for detailed neural correlates is premature.

Langley found that the only processes I proposed were vague and rather standard notions of schema "activation." He also regrets that I did not make more contact with cognitive psychology: For example, to explain the limits on short-term memory or the occasional inaccessibility of long-term memory, or to explain whether separate retrieval mechanisms are used for declarative and procedural knowledge. For the first two questions I have no theory; an attempt at the third will be presented in the forthcoming second edition of Arbib (1972). In any case, the title of the article is "Levels of Modeling of Mechanisms of Visually Guided Behavior"; it was not my charter to apply schemas to other problems in cognition, only to indicate their relevance.

Although there is evidence for some parallelism at the functional level (e.g., in retrieval and classification), **Langley** notes that there is strong evidence for an attentional bottleneck that requires serial processing on tasks such as problem solving and parsing. He notes that recent production-system models (e.g., Rosenbloom & Newell 1987) have addressed this dual character of the human information-processing system; he would welcome a concerted effort to handle these phenomena within a schema-based architecture. In fact, Arbib and Didday (1975) offered a schema-based parallel computational model for the control of serial eye movements – schemas are instantiated to form an assemblage that represents the current environment, and a mechanism like that in Didday's prey-selection model (cf. Section 3.2) uses parallel computation to select the next "focus of attention." More recently, Koch and Ullman (1985) have devised a related model, but based on the internal "spotlight of attention" (Treisman 1983) rather than on overt eye movements. It is of interest that Crick's (1984) model for this phenomenon makes use of the "Malsburg synapses" alluded to in the above section on instantiations.

Schemas and persons. It is the wish of **Gunderson** to understand the relevance of the frog–toad story to the study of human beings. He asks whether human visual-cognitive competencies are to be characterized solely in terms of behaviors and their underlying neurophysiological causes or also in terms of whatever it is like to be aware of making various perceptual discriminations. As in Anderson's accompanying target article [this issue], the use of protocols by Newell and Simon (1972) acknowledges the relevance of introspective data, but a theory of consciousness goes even further, for a nonsentient machine can exhibit measurements of its internal state on a set of dials without in any sense being aware of them. How is the first-person point of view to be "latched onto" conceptually, or is it only a rhetorical starting point that gets lost in the model-building shuffle? For my answer to this, I must refer the reader to Arbib (1985) and Arbib and Hesse (1986) for a schema-based theory of consciousness, not just a theory of overt behavior.

Gunderson asks whether I would adapt Dennett's (1978b) views to construe schemas as hierarchically "homuncularized," with schemas at each level performing diminishing amounts of cognitive donkey-work, thereby accounting for our various intentionalized tasks without spiraling into a regress. The answer is yes, and the "evolutionary rebuttal" to Searle (1980) in Arbib and Hesse (1986) is in the same spirit. Hesse and I explicitly discuss Dennett's views in our Chapter 5, noting the ways in which a theist will find it lacking in offering an analysis of free will. With respect to machine vision, Gunderson asks whether the various levels of computation that allegedly underlie visual cognition are to be viewed as equally psychological. As a crude first approximation, I would answer that low-level vision is not psychological, but that high-level vision is. That this is too crude is suggested by the way the psychological level may influence early processing, as shown, for example, by Weisstein and Wong (1987). He asks how we are to distinguish between processes that are mental even though not introspectively available, and the purely noncognitive physical processes underlying them. I am not sure there is so sharp a dichotomy as the question suggests, but perhaps a start can be made by distinguishing the input/output differences in schemas which can affect our awareness, from the internal workings (e.g., of rhodopsin or synaptic transmitters) which cannot affect it save through their expression in such changed function.

I claim that schemas become "more real" as their functional analysis is refined to allow either a subtler analysis of behavior or an improved mapping of function to neural structure. **Gunderson** suspects that this makes contact with "qualia" or "raw-feels" irrelevant or at most highly covert. A choice of some overall animal behavior as a starting point may be metaphysically theory laden. I agree. In my target article, the focus is neuroethology. When one is "in search of the person" (Arbib 1985), an aspect of person-reality may as well serve as the starting point. A major goal for Mary Hesse and myself in this regard was the reconciliation of personal and social realities. Our approach in Arbib and Hesse (1986) was to analyze the extent to which schemas "within the head" of the individual cohere with social schemas that are distributed across a community and provide part of the external reality for each individual.

Gunderson asks about the mind–body problem. Do I see my schema theory as potentially in the service of a physicalism, but one with a special twist to it (two-way reduction), or as more or less neutral across competing metaphysical accountings of the mind? As described in Arbib and Hesse (1986), schema theory is neutral in that a dualist may see schemas in the mind interacting with, but separate from, the schemas in the brain. However, my own personal position is physicalist: The mind is not separate from the brain but is rather a functional expression of it. Schemas provide the right level of analysis not only for "mind-talk" but for a range of phenomena beyond the range of consciousness (cf. the discussion of Freud in Arbib and Hesse [1986]). We have not yet encountered insuperable barriers to the implementation of these schemas in neural networks. This is not to deny – as the commentaries and this response have made abundantly clear – that the analysis of these schemas and the discovery of their neural underpinnings are still in their early stages, and that we must expect many surprises along the way. It is an exciting time!

NOTE

1. I would like to use this observation to record a hitherto unacknowledged debt. I first coined the term *Rana computatrix* in 1980 or 1981, believing that it "felt right" because I had studied Latin for six years in the 1950s, although I never sat down with a Latin grammar to check this coinage, and harbored the concern that a classicist would one day identify *Rana computatrix* as a Latin malapropism. However, in February 1987 someone (perhaps it was Walter Freeman) observed, "I assume you coined *Rana computatrix* on analogy with W. Grey Walter's *Machina speculatrix.*" I immediately knew that I had, but had never before realized this consciously. A quick trip to the Latin dictionary then revealed that *machina* did indeed mean machine (but also a military engine, and a platform on which slaves were exposed for sale!), but that *speculatrix* meant a (female) observer or watcher – so that *Machina speculatrix* is not, as I had thought, "the speculative machine" but rather "the machine that observes." On this analogy, then, *Rana computatrix* is not "the computational frog" but "the frog that computes," so that, paying due respect to gender agreements, the creature introduced at the end of Professor **Ewert**'s commentary should be *Bufo computator*. However, our debt to Grey Walter (1953) goes beyond etymology, for surely his *M. docilis, M. labyrinthea, M. sopora*, and *M. speculatrix* are the forerunners not only of *Rana computatrix* but also of Braitenberg's vehicles (first

encountered in Braitenberg 1965). In any case, though happily accepting approval of its contribution to synthetic psychology, I still see the need for *Rana computatrix* to evolve as a model for the data of neuroethology (such as those presented by Ewert in his accompanying target article).

References

Aertsen, A. M. H. J. & Johannesma, P. I. M. (1981) The spectro-temporal receptive field: A functional characterization of auditory neurons. *Biological Cybernetics* 42:133–43. [PIMJ]

Amari, S. & Arbib, M. A. (1977) Competition and cooperation in neural nets. In: *Systems neuroscience*, ed. J. Metzler. Academic Press. [arMAA]

an der Heiden, U. (1980) *Analysis of neural networks*. Srpinger-Verlag. [rMAA]

Anderson, J. R. (1983) *The architecture of cognition*. Harvard University Press. [PL]

Antal, M., Matsumoto, N. & Székely, G. (1986) Tectal neurons of the frog: Intracellular recording and labeling with cobalt electrodes. *Journal of Comparative Neurology* 246:238–53. [aMAA, NM, GS]

Arbib, M. A. (1970) On modelling the nervous system. In: *Principles and practice of bionics*, ed. H. E. von Gierke, W. D. Keidel & H. L. Oestereicher. Technivision Services. [rMAA]

(1972) *The metaphorical brain*. Wiley Interscience. [rMAA]

(1975) Artificial intelligence and brain theory: Unities and diversities. *Annals of Biomedical Engineering* 3:238–74. [aMAA]

(1981) Perceptual structures and distributed motor control. In: *Handbook of physiology – the nervous system II. Motor control*, ed. V. B. Brooks. American Physiological Society. [arMAA]

(1982a) *Rana computatrix:* An evolving model of visuomotor coordination in frog and toad. In: *Machine intelligence* 10, ed. J. E. Hayes, D. Michie & Y. H. Pao. Ellis Horwood. [aMAA]

(1982b) Modelling neural mechanisms of visuomotor coordination in frog and toad. In: *Competition and cooperation in neural nets*, ed. S. Amari & M. A. Arbib. *Lecture notes in biomathematics*, vol. 45. Springer-Verlag. [aMAA]

(1982c) Perceptual-motor processes and the neural basis of language. In: *Neural models of language processes*, ed. M. A. Arbib, D. Caplan & J. C. Marshall. Academic Press. [rMAA]

(1985) *In search of the person: Philosophical explorations in cognitive science*. University of Massachusetts Press. [arMAA]

(1987) *Brains, machines, and mathematics*, 2nd ed. Springer-Verlag. [aMAA]

(1987a) Modularity and interaction of brain regions underlying visuo-motor coordination. In: *Modularity in knowledge representation and natural language understanding*, ed. J. L. Garfield. MIT Press/Bradford Books. [rMAA]

Arbib, M. A. & Amari, S. (1985) Sensori-motor transformations in the brain (with a critique of the tensor theory of cerebellum). *Journal of Theoretical Biology* 112:123–55. [aMAA]

Arbib, M. A., Boylls, C. C. & Dev, P. (1974) Neural models of spatial perception and the control of movement. In: *Cybernetics and bionics*, ed. W. D. Keidel, W. Handler & M. Spreng. Oldenbourg. [aMAA]

Arbib, M. A. & Caplan, D. (1979) Neurolinguistics must be computational. *Behavioral and Brain Sciences* 2:449–83. [rMAA]

Arbib, M. A., Caplan, D. & Marshall, J. C., eds. (1982) *Neural models of language processes*. Academic Press. [rMAA]

Arbib, M. A., Conklin, E. J. & Hill, J. C. (1987) *From schema theory to language*. Oxford University Press. [arMAA]

Arbib, M. A. & Didday, R. L. (1975) Eye movements and visual perception: A "two visual systems" model. *International Journal of Man-Machine Studies* 7:547–69. [rMAA]

Arbib, M. A. & Hesse, M. B. (1986) *The construction of reality*. Cambridge University Press. [arMAA]

Arbib, M. A. & House, D. (1987) Depth and detours: An essay on visually guided behavior. In: *Vision, brain, and cooperative computation*, ed. M. A. Arbib & A. R. Hanson. MIT Press/Bradford Books. [aMAA]

Arbib, M. A., Iberall, T. & Lyons, D. (1985) Coordinated control programs for movements of the hand. *Experimental Brain Research Supplement* 10:111–29. [arMAA, HTAW]

Baird, B. (1986a) Nonlinear dynamics of pattern formation and pattern recognition in the rabbit olfactory bulb. In: *Proceedings of the Conference on Evolution, Games and Learning*, ed. D. Farmer, N. Packard, A. Lapedes & B. Wendroff. *Physica* 22D. North-Holland. [rMAA, BB]

(1986b) Bifurcation analysis of oscillating neural network model of pattern recognition in the rabbit olfactory bulb. In: *Neural networks for computing*, ed. John Denker. American Institute of Physics Conference Proceedings 151. [rMAA, BB]

Barlow, H. (1953) Summation and inhibition in the frog's retina. *Journal of Physiology (London)* 119:69–88. [aMAA]

Barlow, H. B., Blakemore, C. & Pettigrew, J. D. (1967) The neural mechanism of binocular depth discrimination. *Journal of Physiology* 193:327–42. [aMAA]

Barnden, J. A. (1985) Diagrammatic short-term information processing by neural mechanisms. *Cognition and Brain Theory* 7:285–328. [rMAA, JAB]

(1986) Complex cognitive information processing: A computational architecture with a connectionist implementation. Technical Report 211, Computer Science Department, Indiana University, Bloomington. [rMAA, JAB]

Bartlett, F. C. (1932) *Remembering*. Cambridge University Press. [aMAA, AKM]

Barto, A. G. (1985) Learning by statistical cooperation of self-interested neuron-like computing elements. *Human Neurobiology* 4:229–56. [rMAA]

Baynes, K., Holtzman, J. D. & Volpe, B. T. (1986) Components of visual attention. *Brain* 109:99–114. [JCM]

Berkinblit, M. B., Feldman, A. G. & Fukson, O. I. (1986) Adaptability of innate motor patterns and motor control mechanisms. *Behavioral and Brain Sciences* 9:585–638. [rMAA]

Berti, A., Papagno, C. & Vallar, G. (1986) Balint syndrome: A case of prosopagnosia. *talian Journal of Neurology* 7:261–64. [JCM]

Bisiach, E. & Luzzatti, C. (1978) Unilateral neglect of representational space. *Cortex* 14:129–33. [JCM]

Bisiach, E., Perani, D., Vallar, G. & Berti, A. (1986) Unilateral neglect: Personal and extra-personal. *Neuropsychologia* 24:759–67. [JCM]

Boylls, C. C. (1975) A theory of cerebellar function with applications to locomotion. I. The physiological role of climbing fiber inputs in anterior lobe operation. Technical Report 75C-6, Computer and Information Science Department, University of Massachusetts at Amherst. [arMAA]

(1976) A theory of cerebellar function with applications to locomotion. II. The relation of anterior lobe climbing fiber function to locomotor behavior in the cat. Technical Report 76-1, Computer and Information Science Department, University of Massachusetts at Amherst. [arMAA]

Brachman, R. (1983) What Is-A is and isn't: An analysis of taxonomic links in semantic networks. *IEEE Computer* 16 (10):30–36. [rMAA; JKT]

(1985) I lied about the trees. *AI Magazine* 6(3):80–93. [rMAA, JKT]

Brachman, R. & Levesque, H., eds. (1985) *Readings in knowledge representation*. Morgan Kaufmann. [JKT]

Brady, M. & Paul, R. (1984) *Robotics research*. MIT Press. [JPW]

Braitenberg, V. (1965) Taxis, kinesis and decussation. In: *Cybernetics of the nervous system*, ed. N. Wiener & J. P. Schadé. *Progress in Brain Research* 17:210–22. [rMAA]

(1984) *Vehicles: Experiments in synthetic psychology*. Bradford Books/MIT Press. [rMAA, DL]

Bullock, T. H. (1983) Implications for neuroethology from comparative neurology. In: *Advances in vertebrate neuroethology*, ed. J.-P. Ewert, R. R. Capranica & D. J. Ingle. Plenum Press. [J-PE]

Burghagen, H. (1979) Der Einfluß von figuralen, visuellen Mustern auf das Beutefangverhalten verschiedener Anuren. Dissertation, University of Kassel. [J-PE]

Caine, H. S. & Gruberg, E. R. (1985) Ablation of nucleus isthmi leads to loss of specific visually elicited behaviors in the frog *Rana Pipiens*. *Neuroscience Letters* 54:307–12. [aMAA]

Cajal, S. R. (1911) *Histology du Système Nerveux de l'homme et des Vertébrés*, vol. 2. Paris: Maloine. [GS]

Cervantes-Perez, F. (1985) Modelling and analysis of neural networks in the visuomotor system of anuran amphibia. Ph.D. dissertation and Technical Report 85-27, Department of Computer and Information Science, University of Massachusetts at Amherst. [arMAA, J-PE, GS]

Cervantes-Perez, F., Lara, R. & Arbib, M. A. (1985) A neural model of interactions subserving prey-predator discrimination and size preference in anuran amphibia. *Journal of Theoretical Biology*. 113:117–52. [arMAA, GS]

Cherniak, C. (1987) Logic and anatomy. Paper read at workshop for the Minnesota Center for Philosophy of Science. University of Minnesota. [KG]

Collett, T. (1977) Stereopsis in toads. *Nature* 267:349–51. [aMAA]

(1982) Do toads plan routes? A study of the detour behaviour of *Bufo viridis*. *Journal of Comparative Physiology* 146:261–71. [aMAA]

Collett, T. & Udin, S. (1983) The role of the toad's nucleus isthmi in prey-catching behaviour. *Proceedings of second workshop on visuomotor* *coordination in frog and toad: Models and experiments*, ed. R. Lara & M. A. Arbib. COINS-Technical Report 83-19, University of Massachusetts at Amherst. [aMAA]

Collett, T. S., Udin, S. B. & Finch, D. J. (in press) A possible mechanism for stereopsis in anurans. [aMAA]

Craik, K.J.W. (1943) *The nature of explanation*. Cambridge University Press. [aMAA]

Crick, F. (1984) Function of the thalamic reticular nucleus: The searchlight hypothesis. *Proceedings of the National Academy of Sciences of the United States of America* 81:4586–90. [rMAA]

Dakin, G. A. & Arbib, M. A. (1986) The hierarchical control of skilled hand movements. COINS Technical Report 86-35, Department of Computer and Information Science, University of Massachusetts at Amherst. [rMAA]

Davis, L. S. & Rosenfeld, A. L. (1981) Cooperating processes for low-level vision. *Artificial Intelligence*, 17:245–63. [aMAA]

Dennett, D. C. (1978a) Why not the whole iguana? *Behavioral and Brain Sciences* 1:103–4. [rMAA, DL]

(1978b) *Brainstorms*. MIT Press/Bradford Books. [rMAA, KG]

De Renzi, E. (1982) *Disorders of space exploration and cognition*. Wiley. [JCM]

Dev, P. (1975) Perception of depth surfaces in random-dot stereograms: A neural model. *International Journal of Man-Machine Studies* 7:511–28. [aMAA]

Dewdney, A. K. (1987) Computer recreations. *Scientific American* March, 256:16–24. [rMAA]

Didday, R. (1976) A Model of visiomotor mechanisms in the frog optic tectum. *Mathematical Biosciences* 30:169–80. [aMAA, JPW]

Didday, R. L. & Arbib, M. A. (1975) Eye movements and visual perception: 'two visual systems' model. *International Journal of Man-Machine Studies* 7:547–69. [aMAA]

Eggermont, J. J., Johannesma, P.I.M. & Aertsen, A.M.H.J. (1983) Reverse correlation methods in auditory research. *Quarterly Review of Biophysics* 16:341–414. [PIMJ]

Epstein, S. (1979) Vermin users manual. Unpublished project report, Department of Computer and Information Science, University of Massachusetts at Amherst. [aMAA]

Evarts, E. V. & Tanji, J. (1976) Reflex and intended responses in motor cortex pyramidal tract neurons in monkey *Journal of Neurophysiology* 39:1069–80. [rMAA]

Ewert, J.-P. (1969) Quantitative Analyse von Reiz-Reaktionsbeziehungen bei visuellem Auslösen der Beutefangwendereaktion der Erdkröte (*Bufo bufo* L.). *Pflügers Archiv* 306:210–18. [J-PE]

(1974) The neural basis of visually guided behavior. *Scientific American* 230:34–42. [aMAA]

(1976) The visual system of the toad: Behavioural and physiological studies on a pattern recognition system. In: *The amphibian visual system*, ed. K. Fite. Academic Press. [aMAA]

(1981) Neural coding of "worms" and "antiworms" in the brain of toads: The question of hardwired and softwired systems. In: *Brain mechanisms of behavior in lower vertebrates*, ed. P. R. Laming. Cambridge University Press. [J-PE]

(1984) Tectal mechanisms that underlie prey-catching and avoidance behavior in toads. In: *Comparative neurology of the optic tectum*, ed. H. Vanegas. Plenum Press. [rMAA, NM]

Ewert, J.-P., Arend, B., Becker, V. & Borchers, H.-W. (1979) Invariants in configurational prey selection by *Bufo bufo* (L.). *Brain, Behavior and Evolution* 16:38–51. [J-PE]

Ewert, J.-P., Borchers, H.-W. & Wietersheim, A. von (1978) Question of prey feature detectors in the toad's *Bufo bufo* (L.) visual system: A correlation analysis. *Journal of Comparative Physiology* 126:43–47. [J-PE]

Ewert, J.-P., Borchers, H.-W. & Wietersheim, A.von (1979) Directional sensitivity, invariance, and variability of tectal T5 neurons in response to moving configurational stimuli in the toad *Bufo bufo* (L.). *Journal of Comparative Physiology* 132:191–201. [J-PE]

Ewert, J.-P., Framing, E. M., Schürg-Pfeiffer, E. & Weerasuriya, A. (submitted) Single neuron activity in the toad's medulla oblongata in response to visual and tactile stimuli: I. A functional approach toward tectal/bulbar spinal circuitry by extraellular methods. [rMAA, J-PE]

Ewert, J. P. & von Seelen, W. (1974) Neurobiologie und System-Theorie eines visuellen Muster-Erkennungsmechanismus bei Kröten. *Kybernetik* 14:167–83. [aMAA, J-PE]

Finkenstädt, T. (1983) Influence of the optic tectum and prosencephalic structures on visually controlled prey-catching and avoidance behaviors in the fire salamander. In: *Advances in vertebrate neuroethology*, ed. J.-P. Ewert, R. R. Capranica & D. J. Ingle. Plenum Press. [J-PE]

Fodor, J. A. (1983) *The modularity of mind*. MIT Press/Bradford Books. [rMAA, JPW]

(1985) Précis of *The Modularity of Mind*. *Behavioral and Brain Sciences* 8:1–42. [rMAA, JPW]

Freeman, W. J. (1975) *Mass action in the nervous system*. Academic Press. [BB]

Freeman, W. J. & Skarda, C. A. (1985) Spatial EEG patterns, non-linear dynamics and perception: The neo-Sherringtonian view. *Brain Research Reviews* 10:147–75. [rMAA]

Fuster, J. M., ed. (1985) The temporal organization of behavior. *Human Neurobiology* 4:Whole Nos. 2 and 3. [rMAA]

Gibson, J. J. (1979) *The ecological approach to visual perception*. Houghton Mifflin. [rMAA, JPW]

Gielen, C.A.M. & van Zuylen, E. J. (1986) Coordination of arm muscles during flexion and supination: Application of the tensor analysis approach. *Neuroscience* 17:527–39. [aMAA]

Gregory, R. L. (1969) On how so little information controls so much behaviour. In: *Towards a theoretical biology, 2: Sketches*, ed. C. H. Waddington. Edinburgh University Press. [aMAA]

Guthrie, E. R. (1935, 1960) *The psychology of learning*. Peter Smith. [JPW]

Hamsher, K. de S. (1978) Stereopsis and the perception of anomalous contours. *Neuropsychologia* 16:453–59. [JCM]

Hanson, A. R. & Riseman, E. M. (1978) VISIONS: A computer system for interpreting scenes. In: *Computer vision systems*, ed. A. R. Hanson & E. M. Riseman. Academic Press. [aMAA]

Havens, W. S. & Mackworth, A. K. (1983) Representing knowledge of the visual world. *Computer* 16:90–96. [AKM]

Hayes, P. (1974) Some problems and non-problems in representation theory. *Proceedings of the Artificial Intelligence and Simulation of Behavior Conference*, Essex University. [rMAA, JKT]

(1979) The logic of frames. In: *Frame conceptions and text understanding*, ed. D. Metzing. Walter de Gruyter. [rMAA, JKT]

Head, H. & Holmes, G. (1911) Sensory disturbances from cerebral lesions. *Brain* 34:102–54. [arMAA, AKM]

Heinzel, H.-G. & Selverston, A. I. (1985) Proctolin modulation of the gastric oscillator in the lobster stomatogastric ganglion. *Society for Neuroscience Abstracts* 11:146.6. [rMAA]

Himstedt, W. (1982) Prey-selection in salamanders. In: *Analysis of visual behavior*, ed. D. J. Ingle, M. A. Goodale & R.J.W. Mansfield, MIT Press. [J-PE]

Hinton, G. E. (1981) Implementing semantic networks in parallel hardware. In: *Parallel models of associative memory*, ed. G. E. Hinton & J. A. Anderson. Erlbaum. [rMAA, JAB]

Hinton, G. E. & Lang, K. J. (1985) Shape recognition and illusory conjunctions. Proceedings of the 9th International Joint Conference on Artificial Intelligence. Morgan Kaufmann. [JAB]

Holland, J. H. (1975) *Adaptation in natural and artificial systems*. University of Michigan Press. [rMAA, BB]

Hopfield, J. & Tank, D. (1986) Computing with neural circuits. *Science* 233:625–33. [rMAA]

Horn, B.K.P. (1985) *Robot vision*. MIT Press. [JPW]

House, D. (1982) The frog/toad depth perception system – A cooperative/competitive model. In: *Proceedings of the workshop on visuomotor coordination in frog and toad: Models and Experiments*, ed. M. A. Arbib. COINS Technical Report 82-16, University of Massachusetts at Amherst. [aMAA]

(1984) Neural models of depth perception in frog and toad. Ph.D. dissertation, Department of Computer and Information Science, University of Massachusetts at Amherst. [aMAA]

Humphreys, G. W. & Riddoch, M. J. (1987) *To see but not to see: A case study of visual agnosia*. Erlbaum. [JCM]

Iberall, T., Bingham, G. & Arbib, M. A. (1986) Opposition space as a structuring concept for the analysis of skilled hand movements. *Experimental Brain Research Series* 15:158–73. [arMAA]

Indurkhya, B. (1986) Constrained semantic transference: A formal theory of metaphors. *Synthese* 68:515–51. [rMAA]

Ingle, D. (1968) Visual releasers of prey catching behaviour in frogs and toads. *Brain, Behavior, and Evolution* 1:500–18. [aMAA]

(1973) Disinhibition of tectal neurons by pretectal lesions in the frog. *Science* 180:422–24. [aMAA]

(1975) Focal attention in the frog: Behavioural and physiological correlates. *Science* 188:1033–35. [aMAA]

(1976) Spatial visions in anurans. In: *The amphibian visual system*, ed. K. Fite. Academic Press. [aMAA]

(1982) Visual mechanisms of optic tectum and pretectum related to stimulus localization in frogs and toads. In: *Advances in vertebrate neuroethology*, ed. J.-P. Ewert, R. R. Capranica & D. J. Ingle. Plenum Press. [aMAA]

Jeannerod, M. (1981) Le controle de l'oeil sur le geste. *La Recherche* 12(120). [aMAA]

Joanette, Y., Brouchon, M., Gauthier, L. & Samson, M. (1986) Pointing with left vs right hand in left visual field neglect. *Neuropsychologia* 24:391–96. [JCM]

Johannesma, P.I.M. (1981) Neural representation of sensory stimuli and sensory interpretation of neural activity: Neural communication and control. *Advances in Physiological Science* 30:103–125. [PIMJ]

Johannesma, P.I.M. & Aertsen, A.M.H.J. (1982) Statistical and dimensional analysis of the neural representation of the acoustic biotope of the frog. *Journal of Medical Systems* 6(4):399–421. [PIMJ]

Johannesma, P.I.M., Aertsen, A.M.H.J., Boogaard, H. van den, Eggermont, J. & Epping, W. (1986) From synchrony to harmony: Ideas on the function of neural assemblies and on the interpretation of neural synchrony. In: *Brain theory*, ed. G. Palm & A. Aertsen. Springer-Verlag. [PIMJ]

Johannesma, P.I.M. & Eggermont, J. (1983) Receptive fields of auditory neurons in the midbrain of the frog as functional elements of acoustic communication. In: *Advances in vertebrate neuroethology*, ed. J.-P. Ewert, R. R. Capranica & D. J. Ingle. Plenum. [PIMJ]

Julesz, B. (1971) *Foundations of cyclopean perception*. University of Chicago Press. [aMAA]

Kandel, E. R. (1978) *A cell-biological approach to learning*. Society for Neuroscience. [aMAA]

Kase, C. S., Troncoso, J. F., Court, J. E., Tapia, J. F. & Mohr, J. P. (1977) Global spatial disorientation: Clinico-pathological correlations. *Journal of the Neurological Sciences* 34:267–78. [JCM]

Kelso, J. A. S. & Scholz, J. P. (1985) Cooperative phenomena in biological motion. In: *Synergetics of complex systems in physics, chemistry and biology*, ed. H. Haken. Springer-Verlag. [rMAA, HTAW]

Kilmer, W. L., McCulloch, W. S. & Blum, J. (1969) A model of the vertebrate central command system. *International Journal for Man-Machine Studies* 1:279–309. [rMAA]

Klahr, D., Langley, P. & Neches, R., eds. (1987) *Production system models of learning and development*. MIT Press. [PL]

Koch, C. & Ullman, S. (1985) Shifts in selective visual attention: Towards the underlying neural circuitry. *Human Neurobiology* 4:219–27. [rMAA]

Kohonen, T. (1984) *Self-organization and associative memory*. Springer-Verlag. [rMAA]

Kugler, P. N. (1986) A morphological perspective on the origin and evolution of movement patterns. In: *Motor development in children: Aspects of coordination and control*, ed. M. G. Wade & H. T. A. Whiting. Martinus Nijhoff. [rMAA, HTAW]

Landis, T., Cummings, J. L., Benson, D. F. & Palmer, E. P. (1986) Loss of topographic familiarity: An environmental agnosia. *Archives of Neurology* 43:132–36. [JCM]

Lara, R. & Arbib, M. A. (1985) A model of the neural mechanisms responsible for pattern recognition and stimulus specific habituation in toads. *Biological Cybernetics* 51:223–37. [aMAA]

Lara, R., Arbib, M. A. & Cromarty, A. S. (1982) The role of the tectal column in facilitation of amphibian prey-catching behaviour: A neural model. *Journal of Neuroscience* 2:521–30. [aMAA]

Lara, R., Carmona, M., Daza, F. & Cruz, A. (1984) A global model of the neural mechanisms responsible for visuomotor coordination in toads. *Journal of Theoretical Biology* 110:587–618. [arMAA]

Lázár, G. (1984) Structure and connections of the frog optic tectum. In: *Comparative neurology of the optic tectum*, ed. H. Vanegas. Plenum Press. [rMAA, GS]

Lázár, G., Tóth, P., Csank, G. & Kicliter, E. (1983) Morphology and location of tectal projection neurons in frogs: A study with HRP and cobalt filling. *Journal of Comparative Neurology* 215:108–20. [aMAA]

Lee, Y. (1986) A neural network model of frog retina: A discrete time-space approach. Technical Report TR-86-219, Department of Computer Science, University of Southern California. [aMAA]

Lettvin, J. Y., Maturana, H., McCulloch, W. S. & Pitts, W. H. (1959) What the frog's eye tells the frog brain. *Proceedings of the Institute of Radio Engineers* 47:1940–51. [aMAA]

Levesque, H. (1986) Knowledge representation and reasoning. *Annual Reviews of Computer Science* 1:255–87. [JKT]

Lieblich, I. & Arbib, M. A. (1982) Multiple representations of space underlying behavior. *Behavioral and Brain Sciences* 5:627–59. [rMAA]

Lloyd, D. (1987) Mental representation from the bottom up. *Synthese* 70:23–78. [rMAA, DL]

Lorenz, K. (1943) Die angeborenen Formen möglicher Erfahrung. *Zeitschrift für Tierpsychologie* 5:235–409. [J-PE]

Luria, A. R. (1959) Disorders of "simultaneous perception" in a case of bilateral occipitoparietal brain injury. *Brain* 82:437–49. [JCM]

(1973) *The working brain*. Penguin Books. [rMAA]

Luthard, G. & Roth, G. (1979) The relationship between stimulus orientation and stimulus movement pattern in the prey catching behaviour of *Salamandra salamandra*. *Copeia* 3:442–47. [J-PE]

Lyons, D. M. (1986) RS: A formal model of distributed computation for sensory-based robot control. Ph.D. dissertation and COINS Technical Report 86-43, Department of Computer and Information Science, University of Massachusetts at Amherst. [arMAA]

Lyons, D. M. & Arbib, M. A. (in press) A formal model of distributed computation for schema-based robot control. [arMAA]

MacKay, D. M. (1966) Cerebral organization and the conscious control of action. In: *Brain and conscious experience*, ed. J. C. Eccles. Springer-Verlag. [aMAA]

Mackworth, A. K. (1983) Constraints, descriptions and domain mappings in computational vision. In: *Physical and biological processing of images*, ed. O. J. Braddick & A. C. Sleigh. Springer-Verlag. [rMAA, AKM]

(in press) Adequacy criteria for visual knowledge representation. In: *Computational processes in human vision: An interdisciplinary perspective*, ed. Z. W. Pylyshyn. Ablex Publishing. [rMAA, AKM]

Marr, D. (1982) *Vision: A computational investigation into the human representation and processing of visual information*. Freeman. [arMAA, KG, JCM, JPW]

Marr, D. & Poggio, T. (1977) Cooperative computation of stereo disparity. *Science* 194:283–87. [aMAA]

(1979) A computational theory of human stereopsis. *Proceedings of the Royal Society (London) Series B* 204:301–28. [aMAA]

Matsumoto, N. & Antal, M. (1984) Physiological and morphological study of prey and predator detecting neurons in the frog optic tectum. *Proceedings of the 6th Annual Meeting of the Japan Society for General and Comparative Physiology* 1:169. [NM]

Matsumoto, N. & Bando, T. (1978) Intracellular recordings of tectal cells of the frog. *Proceedings of the Japan Academy* 54B:386–90. [rMAA, NM]

(1980) Excitatory synaptic potentials and morphological classification of tectal neurons of the frog. *Brain Research* 192:39–48. [rMAA, NM]

Matsumoto, N., Nagano, K. & Li, Qiu-lin (1986) Analyses of synaptic transmission in the optic tectum of the frog by means of simultaneous recording from optic fiber terminals and tectal neurons. *Proceedings of the 8th Annual Meeting of the Japan Society for General and Comparative Physiology* 3:159. [NM]

Matsumoto, N., Schwippert, W. W. & Ewert, J.-P. (1986) Intracellular activity of morphologically identified neurons of the grass frog's optic tectum in response to moving configurational visual stimuli. *Journal of Comparative Physiology* 159:721–39. [arMAA, NM, GS]

Mayhew, J.E.W. & Frisby, J. P. (1980) The computation of binocular edges. *Perception* 9:69–86. [aMAA]

(1981) Towards a computational and psychophysical theory of stereopsis. *Artificial Intelligence* 17:349–85. [aMAA]

McClelland, J. L., Feldman, J., Bower, G. & McDermott, D. (1986) Connectionist models and cognitive science: Goals, directions, and implications. Report of a workshop on connectionism, National Science Foundation. [JAB]

Meijer, O. G. & Roth, K., eds. (1987) *Complex motor behaviour: Control and learning in sports and everyday life*. North-Holland. [HTAW]

Miller, N. (1986) *Dyspraxia and its management*. Croom Helm. [JCM]

Minsky, M. L. (1961) Steps towards artificial intelligence. *Proceedings of the Institute of Radio Engineers* 49:8–30. [aMAA]

(1975) A framework for representing knowledge. In: *The psychology of computer vision*, ed. P. H. Winston. McGraw-Hill. [aMAA, JKT]

Minsky, M. L. & Papert, S. A. (1969) *Perceptrons: An introduction to computational geometry*. MIT Press. [AKM]

Mountcastle, V. B. (1978) An organizing principle for cerebral function: The unit module and the distributed system. In: *The mindful brain*, by G. M. Edelman & V. B. Mountcastle. MIT Press. [rMAA]

Mpitsos, G. J. & Cohan, C. S. (1986) Convergence in a distributed nervous system: Parallel processing and self-organization. *Journal of Neurobiology* 17:517–45. [rMAA]

Mylopoulos, J. & Levesque, H. (1984) An overview of knowledge representation. In: *On conceptual modeling: Perspectives from artificial intelligence, databases and programming languages*, ed. M. Brodie, J. Mylopoulos & J. Schmidt. Springer-Verlag. [JKT]

Nelson, J. I. (1975) Globality and stereoscopic fusion in binocular vision. *Journal of Theoretical Biology* 49:1–88. [aMAA]

Newell, A. (1980) Physical symbol systems. *Cognitive Science* 4:135–83. [rMAA, JPW]

(1982) The knowledge level. *Artificial Intelligence* 18:87–127. [rMAA, JPW]

Newell, A., Shaw, J. C. & Simon, H. A. (1958) Elements of a theory of human problem solving. *Psychological Review* 65:151–66. [rMAA, PL]

Newell, A. & Simon, H. (1963) GPS, a program that simulates human thought. In: *Computers and thought*, ed. E.A. Feigenbaum & J. Feldman. McGraw-Hill. [KG]

(1972) *Human Problem Solving*. Prentice Hall. [rMAA]

Northcutt, R. G. (1986) Strategies of comparison: How do we study brain evolution? *Verhandlungen der Deutschen Zoologischen Gesellschaft* 79:91–103. [J-PE]

Ogden, J. A. (1985) Autotopagnosia: Occurrence in a patient without nominal aphasia and with an intact ability to point to parts of animals and objects. *Brain* 108:1009–22. [JCM]

Oldfield, R. C. & Zangwill, O. L. (1942) Head's concept of the schema and its application in contemporary British psychology. *British Journal of Psychology* 32:267–86; 33:58–64, 113–29. [AKM]

Orgass, B., Poeck, K., Kerschensteiner, M. & Hartje, W. (1972) Visuo-cognitive performance in patients with unilateral hemispheric lesions. *Zeitschrift für Neurologie* 202:177–95. [JCM]

Pellionisz, A. J. (1984) Coordination: A vector-matrix description of transformations of overcomplete CNS coordinates and a tensorial solution using the Moore-Penrose generalized inverse. *Journal theoretical Biology* 110:353–75. [aMAA]

Pellionisz, A. J. & Llinás, R. (1982) Tensor theory of brain function. The cerebellum as a space-time metric. In: *Competition and cooperation in neural nets. Lecture notes in biomathematics*, ed. S. Amari & M. A. Arbib. Springer-Verlag. [aMAA]

Piaget, J. (1971) *Biology and Knowledge*. Edinburgh University Press. [aMAA]

Poggio, G. F. & Poggio, T. (1984) The analysis of stereopsis. *Annual Reviews of Neuroscience* 7:379–412. [aMAA]

Poizner, H., Klima, E. S. & Bellugi, U. (1987) *What the hands reveal about the brain*. MIT Press/Bradford Books [JCM]

Prager, J. M. & Arbib, M. A. (1982) Computing the optic flow: The MATCH algorithm and prediction. *Computer Vision, Graphics and Image Processing* 24:271–304. [rMAA]

Prazdny, K. (1985) Detection of binocular disparities. *Biological Cybernetics* 52:387–95. [aMAA]

(1987) On the coarse-to-fine strategy in stereomatching. *Bulletin of the Psychonomic Society* 25:92–94. [aMAA]

Prechtl, H.F.R. (1986) Prenatal motor development. In: *Motor development in children: Aspects of coordination and control*, ed. M. G. Wade & H.T.A. Whiting. Martinus Nijhoff. [rMAA, HTAW]

Pylyshyn, Z. W. (1980) Computation and cognition: Issues in the foundations of cognitive science. *Behavioral and Brain Sciences* 3:111–69. [rMAA, JPW]

(1984) *Computation and cognition: Towards a foundation of cognitive science*. MIT Press/Bradford Books. [rMAA, JPW]

Qian, X.-S. (1983) About thinking science. *Natural Journal* No. 8. [TYQ]

Ratcliff, G. (1979) Spatial thought, mental rotation and the right cerebral hemisphere. *Neuropsychologia* 17:49–54. [JCM]

Ratcliff, G. & Davies-Jones, G.A.B. (1972) Defective visual localization in focal brain wounds. *Brain* 95:49–60. [JCM]

Ratcliff, G. & Newcombe, F. (1973) Spatial orientation in man: Effects of left, right, and bilateral posterior cerebral lesions. *Journal of Neurology, Neurosurgery, and Psychiatry* 36:448–54. [JCM]

(1982) Object recognition: Some deductions from the clinical evidence. In: *Normality and pathology in cognitive functions*, ed. A. W. Ellis. Academic Press. [JCM]

Reed, E. S. (1984) From action gestalts to direct action. In: *Human motor actions: Bernstein reassessed*, ed. H.T.A. Whiting. North-Holland. [rMAA, HTAW]

Robinson, D. A. (1981) The use of control systems analysis in the neurophysiology of eye movements. *Annual Reviews of Neuroscience* 4:463–503. [aMAA]

Rosenbloom, P. & Newell, A. (1987) Learning by chunking: A production system model of practice. In: *Production system models of learning and development*, ed. D. Klahr, P. Langley & R. Neches. MIT Press. [rMAA, PL]

Rössel, S. (1983) Binocular stereopsis in an insect. *Nature* 302:821–22. [aMAA]

Roth, G. & Jordan, M. (1982) Response characteristics and stratification of tectal neurons in the toad *Bufo bufo* (L.). *Experimental Brain Research* 45:393–98. [J-PE]

Rumelhart, D. E., Hinton, G. E. & Williams, R. J. (1986) Learning internal representations by error propagation. In: *Parallel distributed processing: Explorations in the microstructure of cognition*, ed. D. E. Rumelhart & J. McClelland. MIT Press/Bradford Books. [rMAA]

Rumelhart, D. & McClelland, J., eds. (1986) *Parallel distributed processing: Explorations in the microstructure of cognition*. MIT Press/Bradford Books. [aMAA, DL]

Schank, R. C. (1982) *Dynamic memory: A theory of reminding and learning in computers and people*. Cambridge University Press. [PL]

Schank, R. & Abelson, R. (1977) *Scripts, plans, goals & understanding: An inquiry into human knowledge structures*. Erlbaum. [aMAA, JKT]

Schmidt, R. A. (1975) A schema theory of discrete motor skill learning. *Psychological Review* 82:225–60. [arMAA]

Schwippert, W. W. & Ewert, J.-P. (submitted) Single neuron activity in the toad's medulla oblongata in response to visual and tactile stimuli: II. Identification of neurons by intracellular recording and labeling techniques. [rMAA, J-PE]

Searle, J. R. (1980) Minds, brains, and programs. *Behavioral and Brain Sciences* 3:417–57. [rMAA, KG]

Selverston, A. I. (1980) Are central pattern generators understandable? *Behavioral and Brain Sciences* 3:535–71. [J-PE]

Sperling, G. (1970) Binocular vision: A physical and a neural theory. *American Journal of Psychology* 83:461–534. [aMAA]

Squire, L. R. (1986) Mechanisms of memory. *Science* 232:1612–19. [aMAA]

Steenstrup, M., Arbib, M. A. & Manes, E. G. (1982) Port automata and the algebra of concurrent processes. *Journal of Computer and System Sciences* 27:29–50. [rMAA]

Stein, R. B. (1982) What muscle variable(s) does the nervous system control in limb movements? *Behavioral and Brain Sciences* 5:535–41. [aMAA]

Suga, N. (1984) The extent to which biosonar information is represented in the bat auditory cortex. In: *Dynamic aspects of neocortical function*, ed. G. M. Edelman, W. E. Gall & W. M. Cowan. Wiley. [J-PE]

Székely, G. (1971) The mesencephalic and diencephalic optic centres in the frog. *Vision Research* Suppl. 3:269–79. [rMAA, GS]

Székely, G. & Lázár, G. (1976) Cellular and synaptic architecture of the optic tectum. In: *Frog neurobiology*, ed. R. Llinas & W. Precht. Springer-Verlag. [arMAA, NM, GS]

Szentágothai, J. & Arbib, M. A. (1975) *Conceptual models of neural organization*. MIT Press. [rMAA]

Teuber, H.-L. & Weinstein, S. (1956) Ability to discover hidden figures after cerebral lesions. *A.M.A. Archives of Neurology and Psychiatry* 76:369–79. [JCM]

Thelen, E. (1986) Development of coordinated movement: Implications for early human development. In: *Motor development in children: Aspects of coordination and control*, ed. M. G. Wade & H. T. A. Whiting. Martinus Nijhoff. [rMAA, HTAW]

Tinbergen, N. (1951) *The study of instinct*. Clarendon Press. [J-PE]

Treisman, A. (1983) The role of attention in object perception. In: *Physical and biological processing of images*, ed. O. J. Braddick & A. C. Sleigh. Springer-Verlag. [rMAA]

Tsotsos, J. (1980) A framework for visual motion understanding. Doctoral dissertation, Department of Computer Science, University of Toronto. [JKT]

(1984) Knowledge and the visual process: Content, form and use. *Pattern recognition* 17 (1):13–17. [JKT]

(1985) Knowledge organization and its role in representation and interpretation for time-varying data: The ALVEN system. *Computational Intelligence* 1:16–32. [rMAA, JKT]

(1987) Representational axes and temporal cooperative computation. In: *Vision, Brain, and Cooperative Computation*, ed. M. A. Arbib & A. R. Hanson. MIT Press/Bradford Books. pp. 361–417. [rMAA, JKT]

Tsukahara, N. (1972) The properties of the cerebello-pontine reverberating circuit. *Brain Research* 40:67–71. [aMAA]

von der Malsburg, C. (1981) The correlation theory of brain function. Internal Report 81-2, Department of Neurobiology, Max-Planck Institute for Biophysical Chemistry, Göttingen. [rMAA]

von der Malsburg, C. & Schneider, W. (1986) A neural cocktail-party processor. *Biological Cybernetics* 54:29–40. [rMAA]

von Uexküll, J. (1909) *Umwelt und Innenwelt der Tiere*. Springer-Verlag. [J-PE]

Walter, W. G. (1953) *The living brain*. Duckworth. [rMAA]

Waltz, D. L. (1975) Generating semantic descriptions from drawings of scenes with shadows. In: *The psychology of computer vision* ed. P. H. Winston. McGraw-Hill. [aMAA]

Weisstein, N. & Wong, E. (1987) Figure-ground organization affects the early visual processing of information. In: *Vision, Brain, and Cooperative Computation*, ed. M. A. Arbib & A. R. Hanson. MIT Press/Bradford Books. [rMAA]

Weymouth, T. E. (1986) Using object descriptions in a schema network for machine vision. Ph.D. dissertation and COINS Technical Report 86-24, Department of Computer and Information Science, University of Massachusetts at Amherst. [arMAA]

Whiteley, A. M. & Warrington, E. K. (1978) Selective impairment of topographical memory: A single case study. *Journal of Neurology, Neurosurgery, and Psychiatry* 41:575–78. [JCM]

Whiting, H.T.A., ed. (1984) *Human motor actions: Bernstein reassessed.* North-Holland. [HTAW]

Woods, W. (1975) What's in a link: Foundations for semantic networks. In: *Representation and understanding,* ed. D. Bobrow & A. Collins. Academic Press. [rMAA, JKT]

Young, A. W., ed. (1983) *Functions of the right cerebral hemisphere.* Academic Press. [JCM]

A major new international journal for 1988!

VISUAL NEUROSCIENCE

EDITOR
Katherine V. Fite
University of Massachusetts/Amherst

Bringing together in one forum a broad range of studies reflecting the diversity and originality of contemporary research in basic visual neuroscience, the contributions will utilize neuroanatomical, neurophysiological, neurochemical, neuroimmunological, and behavioral methodologies in addition to computer-assisted theoretical formulations.

The journal's primary emphasis will be on retinal and brain mechanisms underlying visually-guided behaviors and visual perception.

Among the topics of interest are ● Photoreception and transduction ● Retinal anatomy, physiology and neurochemistry ● Retinothalamic pathways and nuclei ● Thalamocortical pathways, visual cortex and telecephalic correlates of visual perception and cognition ● Subcortical visual pathways; oculomotor and visuomotor functions ● Comparative visual system organization and evolutionary perspectives ● Developmental processes, patterns and neuronal specificity ● Theoretical and computational models of vision and visual processes

Published Quarterly
Annual Subscription Rates

US and Canada: $100 (institutions); $50 (individuals); $35 (students) paid in US dollars or Canadian dollar equivalent.

All other countries: £65 (institutions); £35 (individuals); £25 (students) paid in UK sterling.

Cambridge Journals

Cambridge University Press
32 East 57th St., New York, NY 10022, USA; or
The Edinburgh Building, Cambridge CB2 2RU, England