Levels of Modelling of Neural Interactions Underlying Visuomotor Coordination*

Michael A. Arbib

Center for Systems Neuroscience
Computer & Information Science Department
University of Massachusetts
Amherst, Massachusetts 01003

COINS Technical Report 85-36

---

*This paper was supported in part by the National Institutes of Health grant no. NS14971 and National Science Foundation grant no. ECS-8108818.

LEVELS OF MODELLING OF NEURAL INTERACTIONS UNDERLYING VISUOMOTOR COORDINATION *

Michael A. Arbib

Center for Systems Neuroscience
Department of Computer and Information Science
University of Massachusetts
Amherst, MA 01003, U.S.A.

---

* The models reported herein were developed in collaboration with R. Lara, F. Cervantes, A.S. Cromarty and D. House, T. Iberall and D. Lyons with partial support from NIH grant NS14971 and NSF grant ECS-8108818. Preliminary versions of this paper were presented at the Symposium on "Construction and Testing of Models of CNS Action" at the 29th Congress of the International Union of Physiological Sciences, Sydney, Australia, 1 September 1983; at the symposium on "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, St. Louis, Missouri, 3 April 1984, and at the symposium "From Neuron to Thought" held in Firenze, 29 March 1985.

Abstract

Models of the neural interactions underlying visuomotor coordination are constrained by three types of data: ethological, physiological and anatomical. Modelling proceeds both "top-down", seeking to explain behaviour in terms of intermediate-level functional units called "schemas"; and "bottom-up", analyzing the interactions of networks of neurons. We introduce perceptual and motor schemas as units for the functional description of behaviour intermediate between a purely phenomenological description and an account of the detailed neural mechanisms of behaviour. The language of coordinated control programs expresses how such schemas are orchestrated in visually and tactilely guided behaviour. We exemplify how top-down analysis is constrained by both behavioural and lesion data by offering a schema-theoretic analysis of anuran detour behaviour when confronted by barriers between the animal and its prey. We exemplify how bottom-up analysis is constrained by anatomical and physiological data by analyzing a family of models of the interaction of cells in frog tectum during prey-predator discrimination and draw some general conclusions for the fruitful interaction between theory and experiment.

1. Multiple Levels of Analysis

The proliferation of new data about the brain, such as the exquisitely detailed structural anatomy of individual neurons afforded by the HRP method, threatens to overwhelm us with detail. The thesis of this paper is that, if we wish our study of the brain to make contact with ethological analysis of animal behaviour, then we must try to impose order upon the mass of data by providing theoretical analysis at many different levels. We may distinguish the two extremes of *top-down* modelling, which looks at the interaction of functional units called schemas in explaining some overall behaviour of the animal; and *bottom-up* analysis which starts with the detailed mathematical working out of the interaction of individual neurons in explaining network properties. In between, of course, there are many levels of what we might call "middle-out" analysis, in which we seek to relate functions described at one level to mechanisms described at another.

In our top-down work, we can gain much from interaction with workers in artificial intelligence, gaining concepts from their attempts to program computers to carry out certain tasks in machine vision, or to control the behaviour of robots. At the same time, our own evolving study of the brain can help workers in robotics come to understand how highly parallel interactions can be used in controlling the function of their systems. This level of analysis, too, can make contact with studies of ethologists who are concerned with describing basic stimulus patterns and animal behaviours. However, if our concern is to constrain the multitude of possible programs which could underlie some behaviour by data from the neurosciences, then we can turn to lesion studies at a gross level, and neurophysiology and neuroanatomy at a finer level, to begin to suggest how the functional units of our analysis, the schemas, might be instantiated within the brain of the organism under study.

In Sections 2 and 3 we shall present a general methodology for the study of "Perceptual Structures and Distributed Motor Control", and show how a schema-theoretic analysis, informed by the study of robotics, can provide new insights for the analysis of brain mechanisms for the control of hand movements. In the rest of this paper, we shall then turn to studies of frog and toad to further indicate the many different levels of analysis of visuomotor coordination. We shall start with a top-down analysis of depth and detours; and then turn to a more detailed analysis of nets of neurons in tectum and pretectum which subserve prey-predator discrimination.

In all this, it is worth bearing in mind the fiction by Jose Luis Borges (1975) which tells of a College of Cartographers which so prided itself on the accuracy of its maps that each year it produced maps of greater and greater detail until it reached full scale — and Borges remarks that even to this day one can find portions of this ultimate map pegged to the places in the desert that they represent. The point, of course, is that to be useful a map or a model must omit many of the details, the better to allow us to focus upon issues that are essential. Thus it is that, depending upon the task at hand, we will choose different levels of analysis in what follows. This in no way denies that a model which is successful at one level may be improved by later refinement using mechanisms at lower levels. In fact, an initial model at one level may come to be drastically restructured, even recast, as a result of feedback from finer-grain studies. However, it is our contention that to conduct all modelling at the level of greatest detail without reflecting it back to "higher" levels of analysis can only obscure true understanding of the system under study.

2. Perceptual and Motor Schemas.

If we are to describe temporally extended structures of behaviour, we could in principal do this by measuring the positions of all parts of the body at closely spaced points in time, but such measurements are of limited meaning unless there is some way of imposing order upon them — as Kepler did with the ellipse as a way of ordering the vast body of observations of planetary motion. Thus, in studying behaviour we need compact phenomenological descriptions, and we may call these input/output schemas. These provide a rich, yet compact, descriptive language for movement and its timing, but say nothing about how the movement is generated. At the next level, we would try to provide hypotheses for how the behaviour is generated. As we shall elaborate in this section and the next, we do this in terms of ***coordinated control programs*** of internal ***schemas*** to explain the overt behaviour.

The actions of common human behaviour are determined by a far greater knowledge of the environment than is afforded by current sensory stimulation. Our actions are addressed not only to interacting with the environment in some instrumental way, but also to updating our "internal model of the world" [Craik, 1943; Gregory, 1969; MacKay, 1966; Minsky, 1961]. In a new situation we can recognize familiar things in new relationships and use our knowledge of those things and our perception of the relationships to guide our behaviour in that situation. It thus seems reasonable to posit that the internal model of the world must be built of units, each of which roughly corresponds 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.

A diversity of notions of schema have been used in neurology [Head and Holmes, 1911; Fredericks, 1969], in psychology [Bartlett, 1932; Neisser, 1976 ; Oldfield and Zangwill, 1942 and 1943; Piaget, 1971], in artificial intelligence under names such as frames and scripts [Minsky, 1975; Bobrow and Collins, 1975], and in the study of motor-skills [Schmidt, 1975; Schmidt, 1976].

We use the term *perceptual schema* to denote the process whereby the system determines whether a given domain of interaction [Arbib, 1975] is present in the environment. The state of activation of the schema then determines the credibility of the hypothesis that what the schema represents is indeed present whereas other schema parameters represent properties such as size, location, and motion of the perceived object.

One of the most crucial aspects in understanding human behaviour is expressed by our use of words like "plan" and "goal." In the case of an animal such as a frog, it may be appropriate to simply say that a prey-like stimulus will trigger an appetitive response, without finding it appropriate to use such a phrase as "the goal of the animal's behaviour was to eat." In using goal terminology, we tend to suggest that there is a separate representation of the "goal" used in explicating a "plan." When we use the word "plan" we are not only saying that we as observers can project that the animal's behaviour will achieve some desirable future state, but also that the animal has in some sense a representation of that future state and a representation that a particular action will achieve that state. We may represent a plan in terms of the achievement of a goal, and we may see such a goal both in terms of escaping from a present state and in terms of achieving a future state. (A goal will more often be expressed in a condition that is to be satisfied than in terms of the achievement of a specific state itself.)

A "one-way" view of behaviour proceeds through the following stages:

1. Enunciate goal (anticipation of desired state)

2. Perceive situation (STM: Short Term Memory)

3. Plan (using LTM: Long Term Memory)

4. Act — this implies the use of tests contained in the plan

5. Achieve goal

However, we view behaviour in terms of a continuing ***action-perception cycle***, rather than in terms of a discrete stimulus yielding a discrete response (though that is often a useful approximation). An assemblage of activated perceptual schemas provides an extended representation both of an estimate of environmental state based on earlier sensory processing (somewhat like extrapolation, but more subtle), as well as a representation of goals and needs (which may be seen as internal drive variables in some cases, as in "I want X" approximately being seen as "I am in an X wanting state," as in hunger being a state which signals that food is a goal.) New sensory input updates the schema assemblage, and can itself be action-dependent, as in active touch. Physical interactions can change the bodily state (move to comfortable position; eat; drink; mate, etc.). Internal processes can also serve to update the schema assemblage by queries to LTM. The internal state is also updated by knowledge of the state of execution of current plans.

Neisser [1976] has suggested that the subject's exploration of the visual world is directed by *anticipatory schemas*, which he defines as plans for perceptual action as well as readiness for particular kinds of optical structure. The information picked up modifies the perceiver's anticipations of certain kinds of information that, thus modified, direct further exploration and prepare the perceiver for more information. To tell whether or not any coffee is left in a cup, for example, we may reach out and tilt the cup to make the interior visible. We continue tilting the cup until we see coffee or conclude that the cup is empty. Perception is embedded within the organism's ongoing interaction with its environment, for as the organism moves in a complex environment, making, executing, and updating plans, it must stay tuned to its spatial relationship with its immediate environment, anticipating facets of the environment before they come into view.

As action continues, the current plan may continue to be executed simply with tuning or updating of parameters; or, either through some unexpected occurrence, or through completion of the current plan, there may come a process in which that plan has to be changed, some form of replanning or new planning is 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*, in which certain features of the schema assemblage and plan at a particular time are remembered together, from *skill learning* in which the parameters of a plan are updated over time to better tune them to environmental circumstances. 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.

In some cases, it is useful indeed to think of a perceptual schema as constituting an explicit representation of the environment which planning can draw upon. In other cases, perceptual schemas are implicit, embedded within motor schemas and providing the tuning of parameters, or the setting of bounds for those parameters, directly.

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 seek to study programs that are part of the internal state of the system and that can flexibly guide ongoing action in terms of internal goals or drives and external circumstances. Our thesis, then, is that perception of an object (activating appropriate *perceptual schemas*) involves gaining access to routines for interaction with the object (*access to motor schemas*) but does not necessarily involve 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. Therefore perception 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 behaviour of a complex organism is based on these general premises.

1. The action-perception cycle. As the organism moves — making, executing, and updating plans — it must maintain an up-to-date representation of its spatial 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. Activation of perceptual schemas provides access to related motor schemas but does not necessarily entail execution of these schemas. 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 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 within a motor schema.

3. Coordinated Control Programs

We now offer some hypotheses about the nature of coordinated control programs that coordinate the activation of motor schemas. We first review a few elementary concepts from computer science and control theory. The unifying notion is that of internal state, whether it be the state of a control system or the state of execution of a computer program. In either case, we stress that far from being a bundle of stimulus-response pairings, the behaving organism acts in a way that integrates its current sensory stimulation into a complex internal state that is the true determinant of action.

A *program* for an ordinary electronic computer executes one instruction at a time: it transfers inputs to memory, combines pieces of data, controls output devices and, crucially, chooses the next instruction on the basis of a test. Because of these tests the overt behaviour of the program, the temporal sequence of reading of inputs and emission of output values, depends on both the input values and data values already stored internally. The repeated execution of loops until some test is satisfied yields the basic logical property of algorithms: the number of instructions in a program will usually be far less than the number of times instructions are executed when the program is used to process particular data. The program provides the formal description of the process that generates different output sequences on the basis of differing values of inputs and internal parameters. To take a simple example, consider a man walking to the door. Depending on where he starts, his overt behaviour might require three steps or 30, and each such sequence would constitute a program in the fixed-sequence sense of a theatre program. For us, however, it seems more insightful to hypothesize that a single program , in terms of computer science, underlies all these behaviours. In one formalism, for example, we might represent it as

do *advance one step* **until** *door is reached.*

Here we explain all behaviours of the class in terms of a program with one action, *advance one step,* the execution of which is repeated under the control of a single test, *is door reached?* It is this environmental test — coordinating behaviour with position in the environment — that determines the overall time required for the behaviour.

If we regard such a program as a model of human behaviour, we turn our attention from the release of patterns with fixed numbers of steps to the study of ways in which perceptual mechanisms testing "is door reached?" gate motor mechanisms. At a simple level this gating differs little from negative feedback, but when we move to planning behaviour in a complex environment, the loops within loops seem to call for the richer vocabulary that computer science helps provide. The development of a vocabulary of distributed control programs is a major goal of this section. But we first need some concepts from control theory.

A control signal defined by its intended effect may not achieve that effect, either because of disturbances to the system or because of inaccuracy in the controller's knowledge of the controlled system. *Feedback* is then required to compare actual and intended performance so that a compensatory change in the input may be determined. Overcompensation yields instability; undercompensation results in poor adjustment to "noise." Thus not only is feedback necessary, but it must be properly apportioned if the controller is to obtain smooth, coordinated behaviour.

Control theorists use a block diagram to represent a system. Each box represents a subsystem that is continually active whereas the lines linking the boxes illustrate the transfer of data, showing how the output of one system helps determine the input to another. By contrast the boxes of the flow diagrams used by computer scientists to represent programs represent not subsystems but patterns of activation of subsystems. The computer has various subsystems such as memory registers, arithmetic units, and test units. At any time in a computation, certain data are stored in these subsystems, and one box of the flow diagram is activated in the sense that it is used by the computer to determine what tests and operations are to be carried out by the subsystems and how data are to be transferred among them. The lines of the flow diagram then specify how activation is to be transferred from one instruction to another.

Biological control theory usually studies neural circuitry specialised for the control of a specific function, be it the stretch reflex or the vestibuloocular reflex. Yet most behaviour involves complex sequences of coordinated activity of a number of control systems. Arbib [1981] introduced the notion of *coordinated control program* as a combination of control theory and the computer scientist's notion of a program suited to the analysis of the control of movement. These coordinated control programs can control the time-varying interaction of a number of control systems. In the diagrams representing such a program, there are lines representing both transfer of activation and transfer of data.

For the first example of a coordinated control program, consider Figure 1 [Jeannerod, 1981]. The top half shows a collage of handshapes as the hand moves from its initial position to pick up a ball. We note that not only has the hand moved, but the hand has *preshaped* so that when it has almost reached the ball, it is of 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. By 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 may take over in controlling the final grasp. Figure 2 shows a tentative program for this behaviour, adapted from [Arbib, 1981]. In the top half of the figure, we see three perceptual schemas — schemas whose job it is to find information about the environment, rather than to control movement. We have used

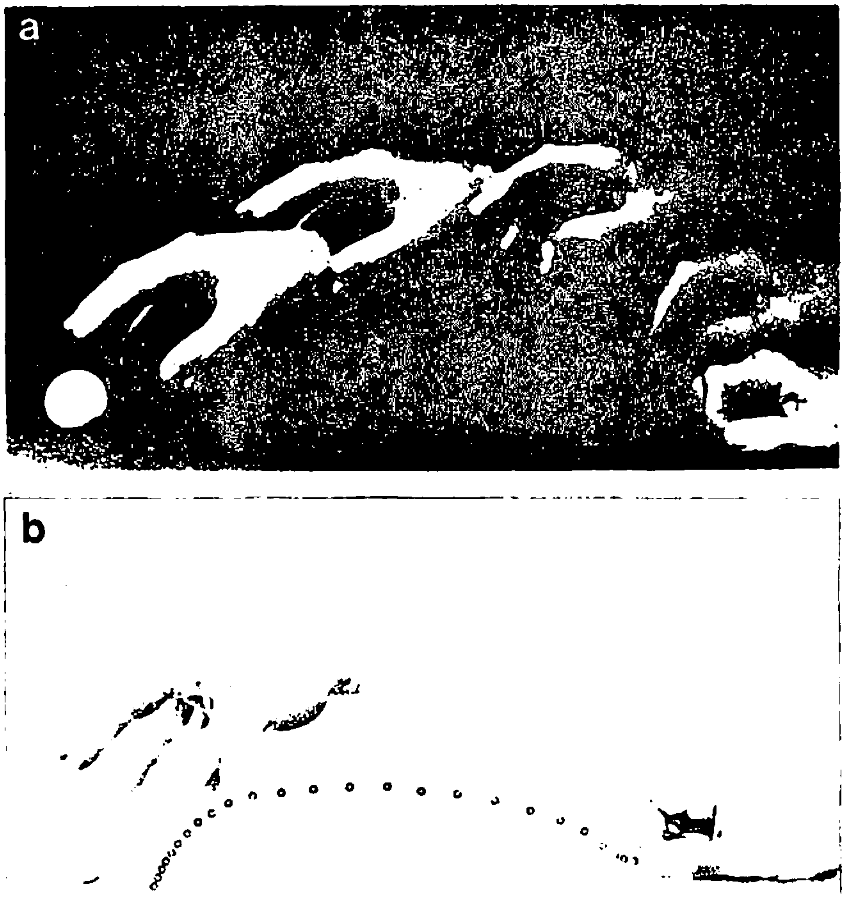

1. (Top) Superimposed view of hand shape and position in reaching for a ball. (Bottom) Circles indicate position of thumbtip in successive movie frames, and indicate separation into a slow phase followed by a fast phase. [Jeannerod 1981]

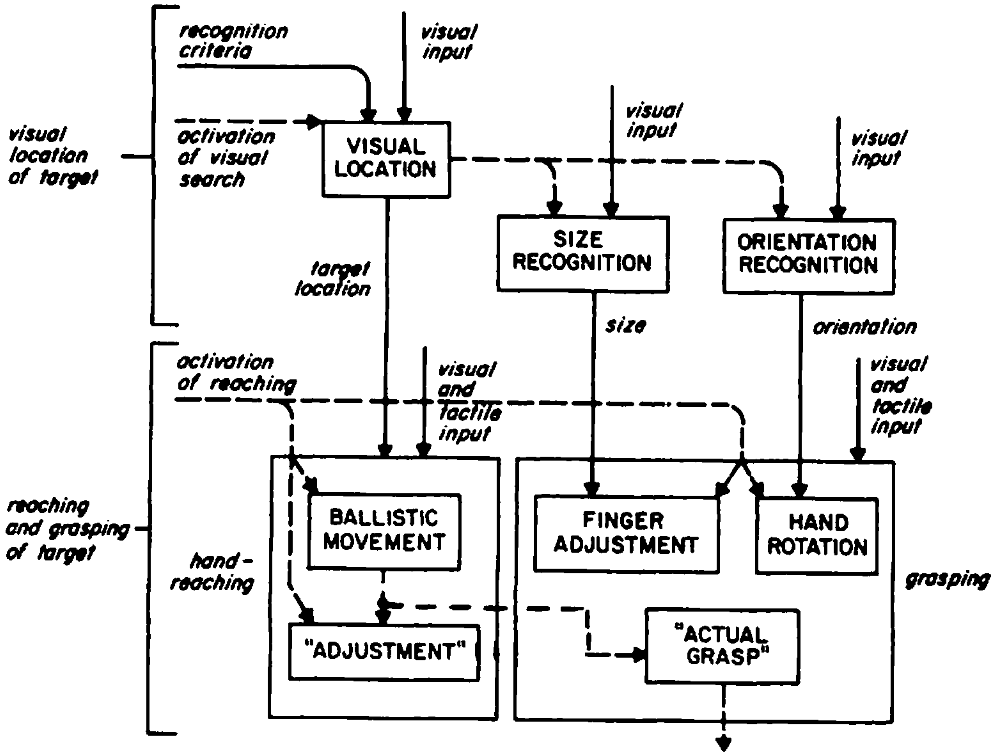

2. A coordinated control program for reaching toward and grasping an object [Adapted from Arbib 1981].

solid lines to indicate the transfer of data from one schema to another; and dashed lines to indicate the transfer of activation. We see that the successful completion of locating the object then activates schemas for recognizing the size and orientation of the object. The output of these perceptual schemas are available for the control of the hand movement. This in turn involves the concurrent activation of two motor schemas, or control systems. One moves the arm to transport the hand towards the object, the other preshapes the hand, with the finger separation and orientation guided by the output of the appropriate perceptual schemas. We note that, once the hand is preshaped, the schema involved in shaping the hand "goes to sleep." It is only the completion of the fast phase of hand transfer that triggers the slow phase of hand transfer as well as "waking up" the final stage of the grasping schema, which will shape the fingers under control of tactile feedback.

With the simple example of reaching to grasp a ball, we have already seen that a coordinated control program involves both concurrency and seriality of schema activation. We have also seen ways in which vision and touch interact. Finally, we have distinguished two types of motor control – a fast ballistic kind, and a slow passage to equilibrium (more of this below). Our next example explores further the possible hierarchical structure of schema programs. Figure 3 shows the task of reaching to grasp a mug. What we stress is that in each of the three cases shown, the movement can be described in terms of three "virtual fingers". In each case, VF 1 (virtual finger 1) is the thumb; VF 2 comprises the 1, 2, or 3 fingers that are slipped through the handle; while VF 3 comprises the remaining fingers, which provide the support beneath the handle to oppose the turning movement of the mug's center of gravity. Arbib, Iberall, and Lyons [1985] hypothesize that the brain controls the movement by using schemas which are structured in terms of three virtual fingers. Thus, preceding the actual movement, a perceptual schema must judge the size of the handle and pass this to a schema which will partition the five actual fingers into the three virtual fingers. Thenceforth, computation proceeds in terms of the three virtual fingers. Following this initial perceptual partitioning, the brain concurrently calls two schemas; one for preshaping the hand, and the other for moving the arm to carry the hand towards the mug handle. If the position and orientation of the wrist is taken as the reference point for an arm movement, then the controller must have knowledge of both the

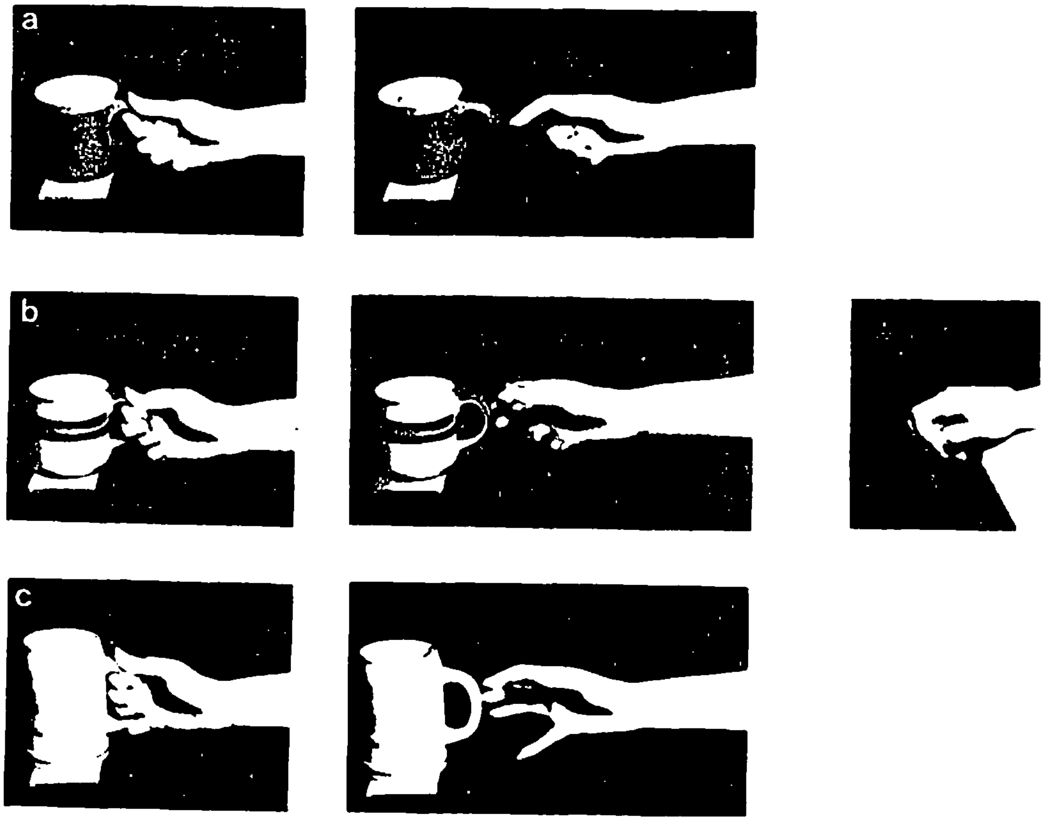

3. Reaching to grasp a mug can be visualized as a task involving three "virtual fingers". [Arbib, Iberall and Lyons 1985]

(internally determined) preshape of the hand and the (obtained from visual perception) position of the handle to compute the appropriate position to which the wrist is to be directed. Once the initial phase has been completed, control is then transferred to hook the VF 2 through the handle and to press VF 1 down upon the position of the handle, as directed by the feel of that handle atop VF 2. We note the subtlety of the positioning of VF 2; it moves in a tangential trajectory around the handle until firmly seated on it. However, to accomplish this movement — which feels as if it is a movement of the fingers along the handle — one must move the wrist in such a way as to allow the fingers to curl around the handle in the appropriate fashion. Thus, the brain must compute an elaborate inverse transformation from tactile information from the fingers to obtain appropriate patterns of muscle contraction for finger, wrist, and arm control.

Figure 4 shows the result of implementing such a program of hierarchical control so that it can drive the graphics display of a hand moving towards a mug. We stress that this program uses, after the initial partitioning schema, schemas in terms of the three virtual fingers. There is no explicit control of the individual fingers as such by the overall coordinated program. The development of such programs is now taking the form of the development of an explicit schema programming language called PRL, the Perceptual Robotics Language. The goal of our research is now three-fold: to use an implementation of this language to control computer graphics to allow us to simulate the effects of different types of program; to use this simulation testbed to come up with the better models of human movement, which can eventually be tied to an understanding of the effects of brain damage, as well as models of the underlying *neural circuitry*; and to study how schema programs may be played out over a network of microprocessors for the distributed dynamic-sensing control of robot behaviour.

We now to look in more detail at the way in which *many movements are divided into an initial fast phase, followed by a slow phase.* We note that in Figure 1 the preshape makes the hand larger than the object to which it is reaching. Analogously, when we place a beer mug down on a table, we can see the movement as divided into a fast phase aimed at a point above the table, and a slow phase which then gently deposits the mug on the

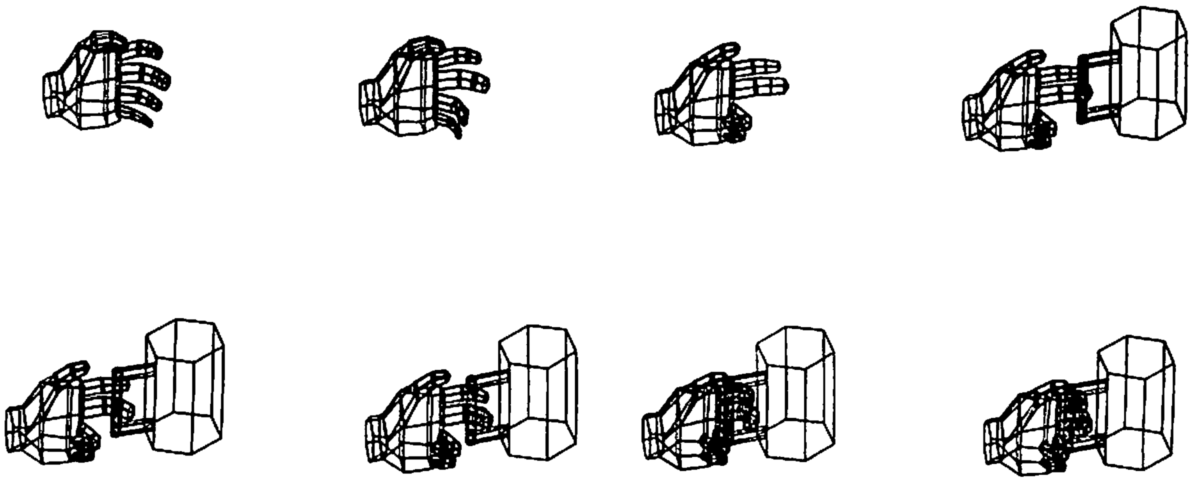

4. Graphics display of output of a simulated coordinated control program structured using the virtual finger concept. [Arbib, Iberall and Lyons 1985]

table. Our suggestion is that, in each case, *the two-phase movement is well structured for a system that is to move quickly, but which is not completely accurate.* The initial form of the preshape, and the initial target of the lowering of the mug are each calculated to provide an "offset" so as to make it highly unlikely that, with any expectable error of movement, contact will be made with the object, respectively with the table, during the fast phase of the movement. However, by so judging this offset that only a small distance remains, the second phase, although executed slowly, will not take much time. This combination achieves speed of overall operation, while avoiding the risk of an unfortunate collision. Arbib, Iberall, and Lyons [1985], have suggested that such a movement may be considered as comprising two calls of a schema MOVE(B,T,C) which will cause a movement toward visually-defined target B, with velocity parameters T, and with an expected tactile termination condition C. We may then regard the overall movement as consisting of two calls to the schema: first, in the mug example (Figure 5), a fast move to the point safely above the table; and secondly, a slow move to make contact with the table. We thus expect normal exit from the first call to be with achievement of the target (within the allowable error) without any contact being made, while the normal exit from the second call should be with achievement of the expected contact. All other cases lead to a failure exit (although we leave open how we are to regard what happens when contact is made at high velocity) and control must then be transferred to some appropriate schema which can handle this situation in a normal way, or carry out the necessary replanning.

While we have suggested that the last program consisted of two consecutive calls of the same schema, but with different parameters, there is evidence that in the case of human movements, two different control systems are involved. It may well be that the initial stage of the movement involves "bang-bang" control, analogous to that involved in making an elevator move from one floor to another in minimal time. Here, the optimal strategy is to maximally accelerate the elevator to the mid-point between the two floors, and then decelerate for the rest of the journey. Clearly, the calculation of such a trajectory requires that the system know both the initial point and the end point of the movement so that the midway transition point can be recognized and acted upon. Similarly, in fast limb movements we can see a burst of initial muscle contraction to move the arm upon its way,

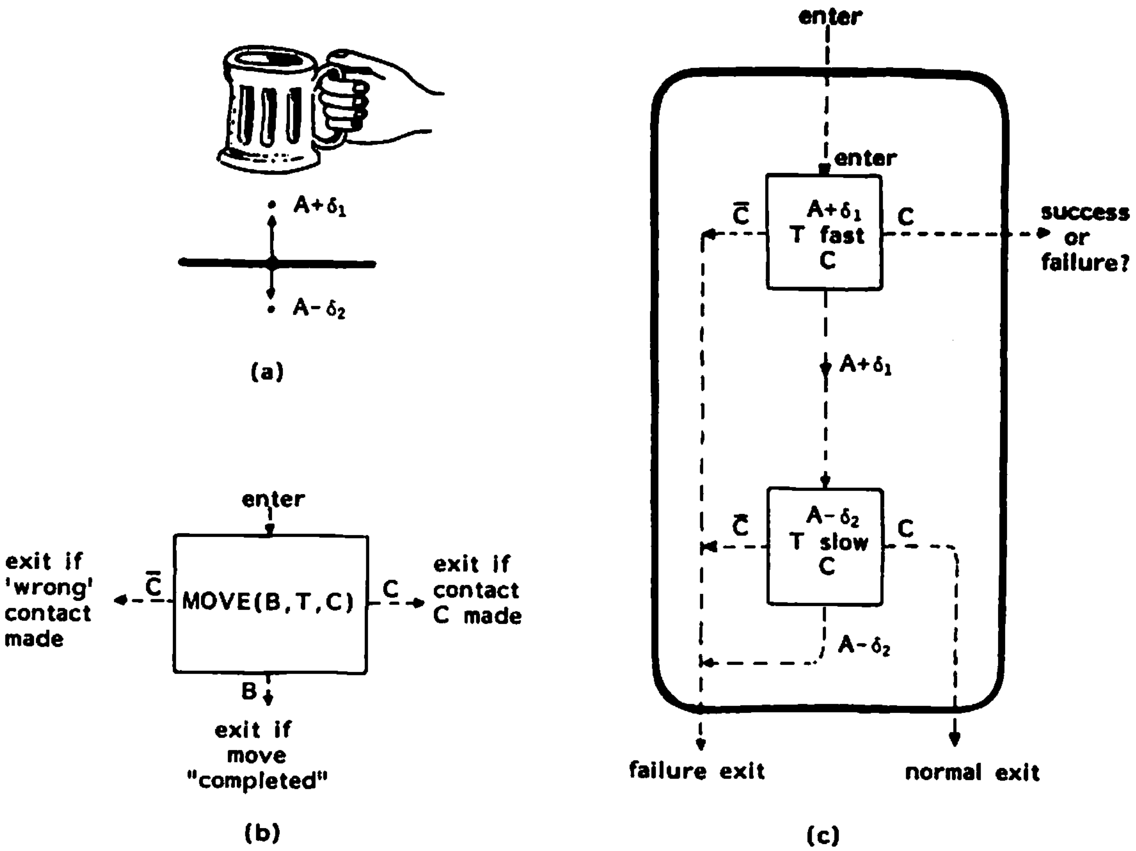

5. (a) The placing of a mug on a table may be seen as composed of two movements: a fast movement to a target $A+\delta_1$ above the intended resting point A on the table (with the "safety undershoot" increasing with an estimate of decreasing accuracy of the movement), and a slow movement towards a point $A-\delta_2$ just below the table, designed to terminate under feedback control on contact. (b) The basic move schema. (c) The coordinated control program for the movement of (a), involving two calls of the schema of (b), each with different parameters. [Arbib, Iberall and Lyons 1985]

and a second burst of antagonist muscle contraction to slow the arm down and bring it to rest at or near the desired point.

However, we may note that the body consists of a skeleton whose movement is controlled by muscles, and that these muscles can be thought of as somewhat like springs whose stiffness is under the control of the nervous system. Thus, we can also bring the body to a desired position by resetting the stiffness, and then letting the muscles relax into a new equilibrium position. However, it must be noted that the desired equilibrium may take quite a while to attain, and may be replaced by a "wrong" intermediate semi-equilibrium if the initial position is too far from the end position. Thus, this mode of control would seem appropriate for the small, slow, second stage of the two stage movement we have described above, rather than for the fast initial stage. We may speak of *feedforward* and *feedback*. A feedback control system is one that moves on the basis of continual sensing of departure from some goal state. A feedforward system uses the sensing of initial state and the representation of the goal to determine a single movement designed to (approximately) achieve that goal.

Possible evidence for this separation of the two stages may be found in observations made by Jeannerod, Michel and Prablanc [1985] on an unfortunate patient who underwent neurosurgery in Lyon. At the end of the surgery, a small piece of surgical gauze was left in one of the arteries. It blocked blood flow to the somatosensory cortex on the left-hand side of the brain. As a result, the patient not only lost the sense of touch for the right-hand side of her body (most acutely in the extremities) but also lost position sense, the ability to sense how the body or parts of it are being held. Figure 6 presents a sequence of pictures showing her moving to grasp a ball. The dashed line separates the visual field, to the left, from the area that she could not see, to the right. Initially, her hand was placed in a bizarre posture by the clinician — and her brain had no way to tell what that shape was. She was able to see the ball and use this information to direct the arm to carry the hand; but lacking the precision information about the initial shape of the hand, she did nothing to start preshaping the hand until the hand came into view. She then preshaped it moderately well, but since only visual shaping, and neither

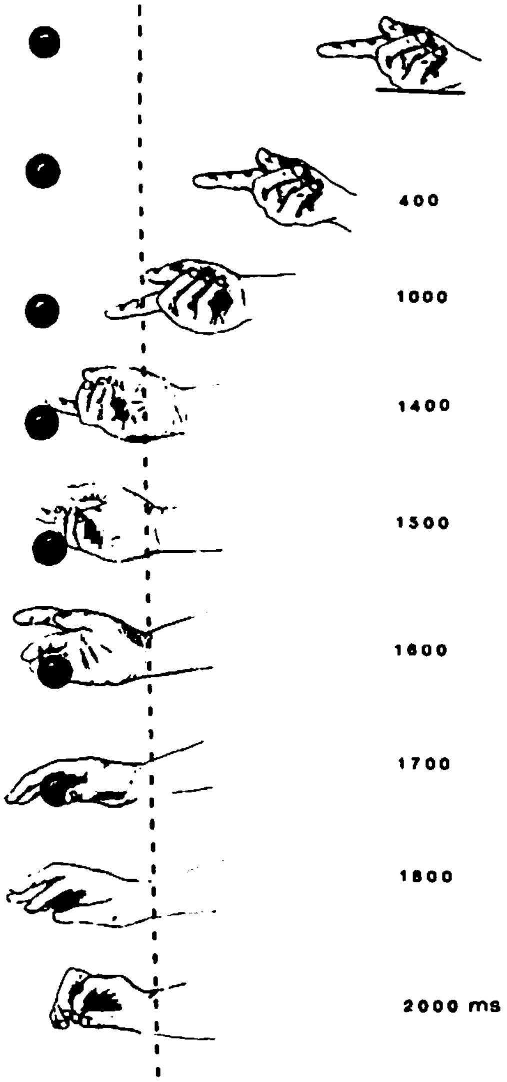

6. Reaching for an object by a patient with lesion of contralateral somatosensory cortex. [Jeannerod, Michel and Prablanc 1984]

tactile information nor position sense were available, the preshaping was somewhat clumsy. Nonetheless, the crude preshape was formed, and then the hand shaped itself around the ball. Note, however, that this grasping is also clumsy, since it is only the mechanical force stopping the movement of the fingers that now controls this movement in addition to visual information — there is no tactile feedback that can be used by the brain. The inability of the patient to preshape her hand prior to its coming into view suggests very strongly that the initial preshape movement is of the ballistic kind in which knowledge of an initial position as well as final position is required to compute the necessary control forces.

4. An Overview of Modelling Problems

Lettvin, Maturana, McCulloch and Pitts [1959] initiated the behaviourally-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. In this spirit, we view the analysis of interactions between layers of neurons as a major approach to modelling "the style of the brain". In Section 5, we shall offer a general view of cooperative computation between neurons within a layer, and between layers within the brain. (The relation of "maps as control surfaces" to the general study of perceptual structures and distributed motor control is given in Arbib [1981].) In following sections, we shall then exemplify these general principles in specific models of cooperative computation in neural circuitry underlying visuomotor coordination in frog and toad. But we first devote ourselves to an explicit discussion of the various grains at which neural networks can be modelled.

We have seen that we may gain understanding into the mechanisms underlying behaviour at the level of schemas, and at the more detailed level of interacting modules or layers. We now want to see what can be learned by analysis at the level of single neurons. We would note that for many behaviours, this more refined level of analysis may be superfluous. However, if our concern is to relate analysis of behaviour to neurophysiological recording or neuroanatomical data, then this further level of analysis is necessary. We must then confront the question of how detailed the analysis of each neuron must be. Hodgkin and Huxley have shown us how much can be learnt from an electrical analysis of membrane properties in understanding the propagation of electrical activity along the axon; and Rall is one of the leaders in showing how the study of membrane properties in dendrite, soma and axon can help us understand form-function relations in the single neuron and in small neural circuits. However, the complexity of such analysis makes it necessary to use a more lumped representation of the individual neuron if we are to understand the properties of large networks.

We may determine units in the brain physiologically — for example, by electrical recording — and anatomically — e.g. by staining. In many regions of the brain, we have an excellent correlation between physiological and anatomical units — we know which anatomical entity yields which physiological response. Unfortunately, this is not yet the case in many studies of visuomotor coordination in frog and toad. We have data on the electrophysiological correlates of animal behaviour, and we have anatomical data. Often, though, we do not know which specific cell, defined anatomically, yields an observed electrophysiological response. For example, we have the Golgi anatomy of the frog tectum, shown in Figure 7a, and the physiological responses recorded from tectum during facilitation of prey-catching behaviour shown in Figure 17d. However, our identification of the physiological responses with specific anatomically defined cells is still hypothetical. Nonetheless, such choices have to be made in formulating and testing our models.

This point can be reinforced by looking at the tectum, which has been the central structure to date in our neural modelling of anuran visuomotor coordination. Figure 7a shows the Golgi study of the frog tectum offered by Szekely and Lazar (1976). We are already struck by the diversity of the neurons, and the richness of their structure. Yet we may note how little the diagram tells us about the detailed connectivity of the neurons — which must surely be essential in relating structure to function. However, there can be such a thing as too much detail, as can be seen from the electron-micrograph of Figure 7b. Whatever the importance of a detailed analysis of such local circuitry, it is clear that a computer simulation of the response of the whole tectum which analyzed each component

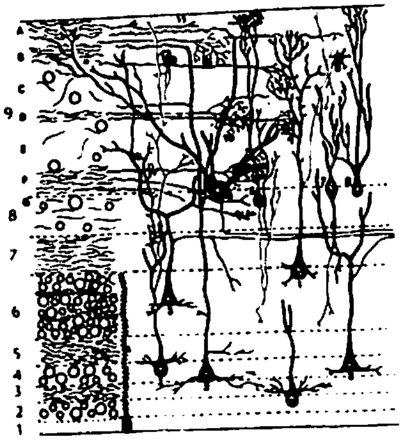
A
B
C
9 D
E
F
G
8
7
6
5
4
3
2
1

(a)

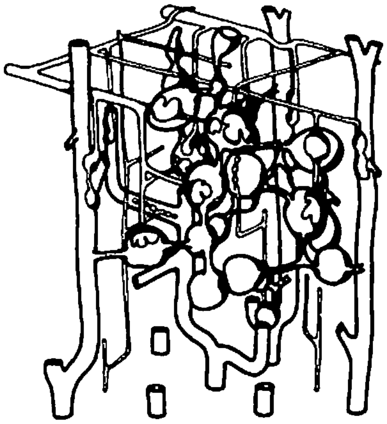

(b)

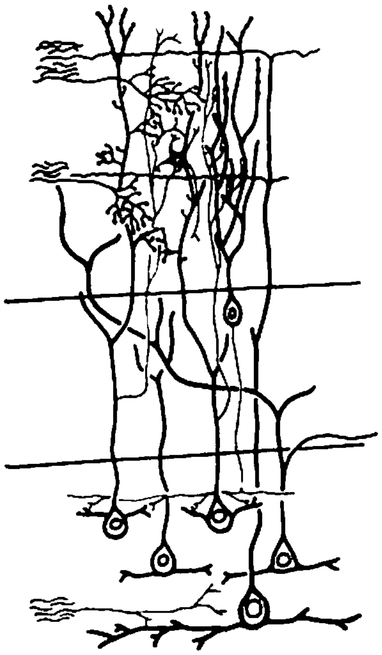

(c)

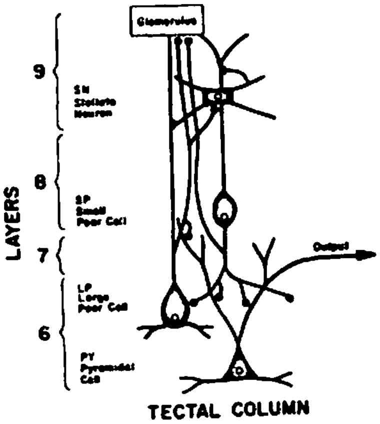
Glomerulus
LAYERS
9
SN
Stellate
Neuron
8
SP
Small
Pear Cell
7
Output
LP
Large
Pear Cell
6
PY
Pyramidal
Cell
TECTAL COLUMN

(d)

7. (a) Diagramatic 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 fibres [from Szekely and Lazar 1976].

(b) Details of synaptic interaction of dendritic appendages, which exceed current models in intricacy [from Szekely and Lazar 1976].

(c) Szekely and Lazar's schematic for a tectal column [from Szekely and Lazar (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 fibres and it inhibits the LP and SP cells. The PY is activated by the LP, SP, and optic fibres, and is the efferent neuron of the tectum. [Arbib 1982b] at this level of detail would be too large to run efficiently on even the most super of the supercomputers. In seeking our own preliminary level of lumping the circuitry of the tectum we were guided by Szekely and Lazar's own figure, Figure 7c, of the "tectal column". This terminology does not imply, as it does in analysis of mammalian cortex, that adjacent columns must have dramatically different response properties. Rather, the suggestion is simply that we have here a basic unit of vertical organization, a "tile" that can be repeated again and again to approximate the tectal neural net. Lara, Arbib and Cromarty (1982) adapted this tectal column in the form shown in Figure 7d, representing each component by a simple differential equation, and used it to model the phenomenon of facilitation: the fact, observed by Ingle (1968), that a frog will not respond to a single presentation of a prey for too short a period, but will respond to a subthreshold length of presentation if it follows within two or three seconds an earlier presentation of such length. In Section 7, we shall outline how an array of such tectal columns has been used in modelling prey-predator discrimination.

In comparing the Golgi anatomy of Figure 7a with the model of Figure 7d, we see that a number of choices have been made. In Figure 7a we see that there are two types of output cells for the tectum, the pyramidal cell and the large tectal ganglionic neuron. Our model assumes that it is only the output of the former that is relevant to the phenomena that we are considering. Clearly, our models must be of such a kind that they are adaptable when we come to phenomena that in fact can be shown to depend upon the ganglionic output. Note, too, that we have ignored the bipolar neurons and amacrine cells, and that we have made certain assumptions about the connectivity between the neurons that are included in the model. However, an important point of our modelling methodology will be that we set up our simulation in such a way that we can use different connectivity on different simulations. In this fashion, we can generate hypotheses which can then be subjected to further experimental test.

Even if we have made a satisfactory choice of how to correlate physiological units with anatomical units, and of the appropriate connectivity, we still have the problem of correlating cellular structure and function with the animal's overall behaviour. In the modelling to be described in Section 7, for example, we have assumed that the activity of the pyramidal cells correlates with the orienting response of the animal. We have also assumed that when a population of pyramidal cells is active, the resultant orientation is to the mean position corresponding to that population, though there is not yet evidence to discountenance further hypotheses, such as that the orientation will be to the spatial locus corresponding to the peak of the pyramidal cell activity.

We have already spoken of the need to have a family of models which allows one to experiment with a number of different connectivities and parameter settings for the cells of the model. There still remains the question of what is the appropriately detailed model. Is it the fact that the overall behaviour of a large collection of cells depends critically on the fine details of the response performance of each individual neuron, or can we hope to use relatively simple, computationally efficient, neuron models and still derive significant information about the behaviour of the population? In the models to be described below, we have described the behaviour of the neuron by a simple differential equation linear in terms of the synaptically weighted input values, and have assumed that the input from one cell to another is given by a simple non-linear transformation of the membrane potential of the source cell. We believe that with such models we can probe whether the neural networks of different kinds can yield overall classes of behaviour. Future research will be both less detailed — trying to provide quantitative analyses correlating classes of neural networks with classes of behaviours–and more detailed, as we try to establish detailed parametric specifications which can be subjected to experimental test in the laboratory. Section 5 will provide a survey of some of the conceptual models that enter into our search for "the style of the brain," while Section 7 will present several stages of our attempts to model certain experimental studies of visuomotor coordination in frog and toad at the level of interacting tectal columns.

Our modelling methodology must be based not on a single "take it or leave it" model, but rather on the exploration of a variety of different connectivities within some overall paradigm of brain function. Thus, the models to be described below are dominated by two main considerations: the visual system of the animal must be considered in the context of the ongoing behaviour of the animal — thus the stress on visuomotor coordination, rather than on vision per se; and the analysis will be in terms of the interaction between concurrently active regions of the brain, rather than in terms of any simple one-way flow of information in a hierarchically organized system. We use the term ***cooperative computation*** to refer to this style of concurrent neural processing.

In addition to our concern for embedding the brain within the ongoing cycle of the animal's action and perception, and studying the brain itself in terms of the cooperative computation of interacting subsystems, the models to be exhibited in Section 7 exhibit a style of "evolutionary" modelling. In Section 7a we present a model of local circuitry in the tectum (a 'tectal column') to explain certain facilitation effects in prey-catching behaviour; then in Section 7b we study a linear array of such columns to model certain data on size-dependence of prey-catching activity in toads; while in Section 7c, we add inhibition from pretectum to such an array to model the behaviour of an animal confronted with more than one prey-stimulus. These models form three stages in an evolutionary sequence for *Rana Computatrix*, our developing model of the neural circuitry underlying visuomotor coordination in frog and toad. Tectum and pretectum are but two of the many brain regions to be incorporated into the model during its further evolution. As a first approximation, we continually try to localise the neural processes underlying some overt behaviour of the animal within some relatively small portion of the brain. As we come to analyze more functions, though, we find that each function may require activity in many portions of the brain, and that each portion of the brain will be involved in many different activities. Thus, having successfully modelled several phenomena, one should try as far as possible, when modelling a new phenomenon, to do it by minor adaptations of the previous model, preserving the earlier successes, rather than introducing an *ad hoc* model of a new brain region specifically designed to achieve the new specified task. In this connection we stress an important property of such models as that of the tectal column. It is not a simple "take it or leave it" model, but rather a family of models. Within the overall structure shown in Figure 7d, there are choices as to the connectivity of various retinal pathways and other projections to the tectum; and there are certainly many choices about the connectvity weightings amongst different neural classes. Through computer simulation or mathematical analysis, one can see how the overall function of the system varies as a property of these choices. Thus, by seeing which choices best match known properties of the system, we can make predictions about other properties which have not yet been measured. In addition, the model then has the flexibility to be altered to help us make sense of changing experimental evidence.

5. Background Models: In Search of the Style of the Brain

Before turning, in Sections 6 and 7, to various stages in the evolution of *Rana Computatrix*, we devote the present section to a number of background models which establish the "style of the brain" with which we approach our modelling of visuomotor coordination in frog and toad.

Since we are concerned with motor control, we of course make use of such concepts as feedback and feedforward. In many treatments of these concepts in the literature on biological control systems, we see the use of lumped models. For example, the direction in which the animal should turn is encoded by a single angle variable. However, since we shall be concerned with the way in which patterns on the retina inpinge upon ongoing activity within the brain, we shall not consider it permissible to regard this angle as explicitly available in the brain as the value of, for example, firing of some neuron. Rather, we must consider it as encoded by the locus of the peak of activity within a neural array. Perhaps the first model of distributed motor control of this kind is that of Pitts and McCulloch (1947).

5a. *Distributed Motor Control.*

Apter (1945, 1946) had shown that each half of the visual field of the cat maps topographically upon the contralateral superior colliculus. In addition to investigating this sensory map, she studied the motor map by strychninizing a single point on the collicular surface, flashing a diffuse light on the retina, and then observing which point in the visual field was affixed by the resultant change in gaze. She found that these sensory and motor maps were almost identical, and this basic finding has been replicated and extended in many recent studies. Starting from these data, Pitts and McCulloch developed their model [Pitts and McCulloch 1947] of the reflex arc that extended from the eyes through the superior colliculus to the ocular-motor nuclei, thereby controlling the muscles that direct the gaze so as to bring the fixation point to the center of gravity of distribution of the visual input's brightness. (Our current knowledge of retinal preprocessing enables us to substitute for the term *brightness* such a term as *contour information* or an expression that describes some other feature of the input.) Pitts and McCulloch noted that excitation at a point on the left colliculus corresponds to excitation from the right half of the visual field and so should induce movement of the eye to the right; gaze is centered when excitation from the left is exactly balanced by excitation from the right. Their model is so arranged that each motor neuron controlling muscle fibers in the muscles that contract to move the eyeballs to the right, for example, should receive excitation summing the level of activity in a thin transverse strip of the left colliculus. This process provides all the excitation for muscles turning the eye to the right. Reciprocal inhibition by axons from nuclei of the antagonist eye muscles, which are excited similarly by the other colliculus, performs subtraction. The quasi-center of gravity's vertical coordinate is computed similarly. Eye movement ceases when and only when the fixation point is the center of gravity. Such a model leads to the idea that a plausible subsystem for vertebrate nervous systems may be one in which position of the input on the control surface encodes the target to which the muscular control will be sent. Of course, much remains to be done in turning such a general scheme for distributed motor control into a specific model of a specific system. For example, the Pitts-McCulloch model does not give an account of ballistic movements. Again, it does not show us how, for increasing angles of deviation of the target, visual tracking might first evoke movement of eyes alone, then of eyes and head, and then of eyes, head, and trunk. It remains an important task in brain theory to explain how the output of a motor computer would control not a single pair of antagonist muscles, but rather a whole hierarchy of subcontrollers, in a distributed way.

5b. *A Model of Frog's Snapping.*

Another problem is that in much visually guided behaviour, the animal does not simply respond to "the center of gravity" of visual stimulation, but rather is responding to some property of the overall configuration. Consider, for example, the snapping behaviour of frogs confronted with one or more fly-like stimuli.

Ingle (1968) found that in a certain region around the head of a frog, the presence of a fly-like stimulus elicits a snap; that is, the frog turns so that its midline is pointed at the stimulus and saps it with its tongue. When confronted with two "flies," 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 (1970, 1976) offered the simple model of this choice behaviour shown in Figure 8a. It is presented not as the state of the art — in fact, we shall see a more recent model built upon it in Section 7 – but rather as a clear example of the processing of structured stimuli to provide the input to a distributed motor controller akin to that discussed in Section 5a. 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, with new circuitry introduced whereby different regions of the tectum compete in such a way that in normal circumstances only the most active region provides an above-threshold input to the motor circuitry. To achieve this effect we first introduce 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 employs 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

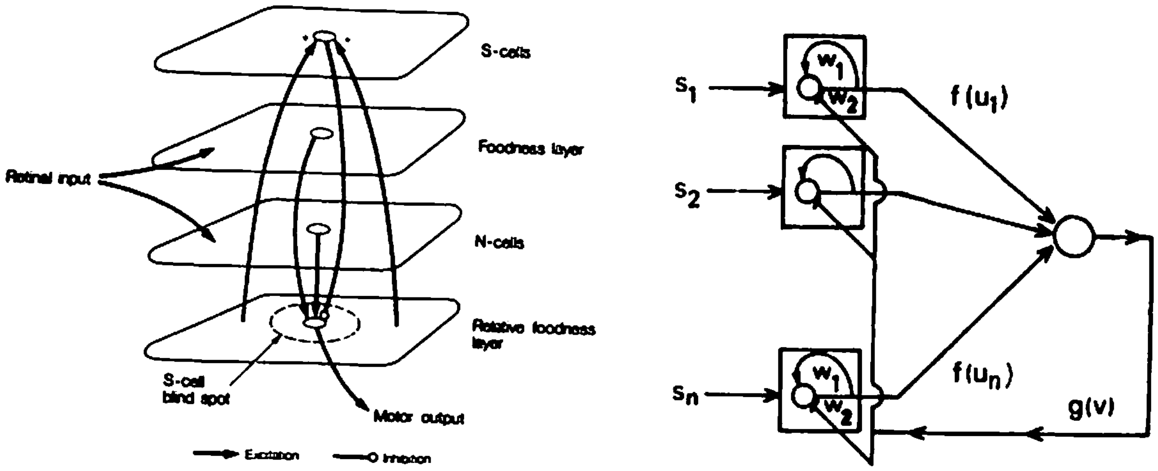

(a)

(b)

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 and Arbib 1977].

the corresponding cells in the foodness layer by an amount that augments with increasing 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.

When we examine the behaviour 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 close activity 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 are 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 to suddenly become more active, then the deadlock should 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 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. Didday hypothesized that the S-cells and N-cells modelled the "sameness" and "newness" cells, respectively, that had been observed in the frog tectum.

Regrettably, no experiments have been done to test this hypothesis.

5c. *Competition and Cooperation in Neural Nets.*

The above model of prey selection is an example of a broad class of models dealing with competition and cooperation in neural nets. As one example of a model of such a kind, let us consider the problem of stereopsis, or segmentation on depth cues. Julesz (1971) has designed "random-dot stereograms" in which each eye receives a totally random pattern, but in which there are correlations between the inputs to the two eyes. Specifically, the different regions in the two inputs are identical save for a shift in position, yielding a different disparity in the two retina (Figure 9a). Although such a pattern for a naive subject can initially appear to be nothing but visual noise, eventually disparity matching takes place and the subject perceives surfaces at different depths. Barlow, Blakemore and Pettigrew (1967) and Pettigrew, Nikara, and Bishop (1968) have found that cells in cat visual cortex are tuned for retinal disparity, and similar cells are posited in the human. What presumably causes the initial perception of visual noise is that in addition to the correct correlation of points in the two retinas, there are many spurious correlations, and computation is required to reduce them (Figure 9b).

Dev (1975) [see also Sperling (1970), Arbib, Boylls and Dev (1974), Nelson (1975), and Marr and Poggio (1977)] has proposed that the cells of a given disparity be imagined as forming a population arrayed in a spatial map corresponding to the map of visual direction. Connections between cells could then be arranged so that nearby cells of a given disparity would be mutually excitatory, whereas cells nearby in visual direction but different in disparity would have inhibitory interaction (Figure 9c). In this way, the activity of the array would organize into a pattern where in each region of visual direction, cells of only one disparity type would be highly active. As a result the visual input would eventually be segmented into a number of distinct surfaces.

In the stereopsis model, then, we have competition in the disparity dimension and cooperation in the other dimensions. The Didday model (Figure 8a) can be regarded

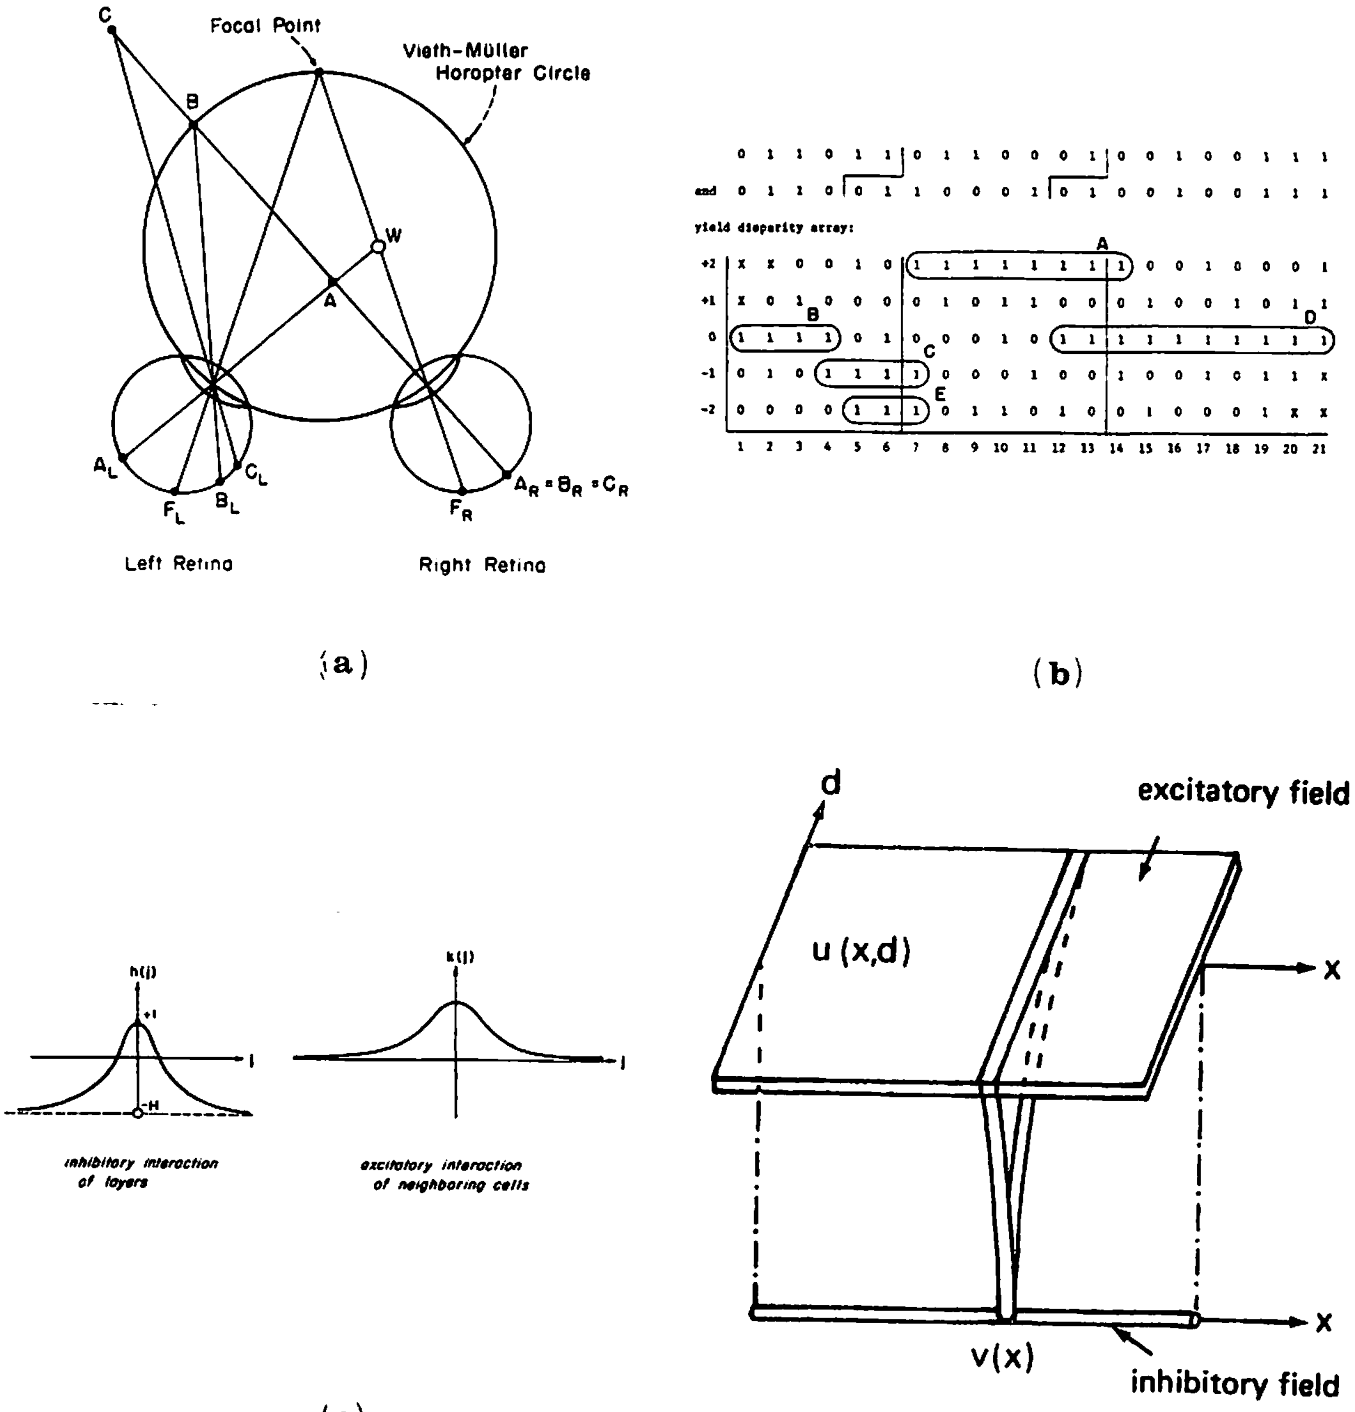

(a)

(b)

(c)

(d)

9. (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 difference in disparity increases. [From Arbib, Boylls and Dev (1974).] (d) The full model of competition and cooperation which allows the idea shown in (c) to be subject to mathematical analysis [from Amari and Arbib 1977].

as the limiting case where there is only a competition dimension, namely that of prey location. Such informal observations have laid the basis for rigorous mathematical analysis of competition and cooperation in neural nets. For example, Amari and Arbib (1977) both offer the "primitive cooperation model" of Figure 8b which allows us to gain a mathematical handle on Didday's results, as well as a more sophisticated model, shown in Figure 9d, which allows us to provide a stability analysis of a model of the kind studied by Dev for stereopsis. Amari (1982) gives an up-to-date perspective on such models.

6. Depth and Detours

As is clear from the models treated in the previous section, a 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 spatio-temporal function without worrying about its quantization in terms of the activity of the separable neurons. We now present a second model of depth perception of this kind, and then use it to argue the need for high-level modelling in the design of neuroethological experiments.

We have seen that the problem for many models of binocular perception is to suppress ghost targets, and that the essence of the Dev scheme was to have those neurons which represent similar features at nearby visual directions and approximately equal depths excite each other, whereas those neurons which correspond to the same visual direction but different depths were (via interneurons) mutually inhibitory. In this way, neurons which could represent elements of a surface in space will cooperate, whereas those which would represent paradoxical surfaces at the same depth will compete. The result is that, in many cases, the system will converge to an adequate depth segmentation of the image.

However, as we shall now see, such a system may need extra cues. For example, in looking at a paling fence, if several fenceposts are matched with their neighbors on the other eye in a systematic fashion, then the cooperative effect can swamp out the correct pairing and lead to the perception of the fence at an incorrect depth. In animals with frontal facing eyes such ambiguity can be reduced by the use of vergence information to drive the system with an initial depth estimate. Another method, as House posits, is to use accommodation information to provide the initial bias for a depth perception system; this is more appropriate to the amphibian, with its lateral-facing eyes.

Ingle (1976) had observed that a monocular frog can snap fairly accurately at prey presented within its monocular field, implying that depth can be estimated from monocular cues. Collett (1977) used experiments with prisms and lenses placed in front of the eyes of the toad to show that, in its binocular field, the toad relied mainly on stereopsis, but that the monocular toad did make depth judgments based on accommodation. The problem, then, was to design a model which would function on accommodation cues in the monocular animal but which would nonetheless be most dependent upon binocular cues within the binocular field. The resultant model (House, 1982) is shown in Figure 10. 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, M and S are so intercoupled that a point in the accommodation field will excite the corresponding point in the accommodation field, 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 11 shows stages in the processing by this model of a scene comprising a fence and two worms. The left-hand column of each figure 11b shows the accommodation (above) and the disparity field (below) for the fence information. In the top image of Figure 11b, 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 intial state of the stereopsis field. The targets are better tuned, but they offer ghost images in addition to the correct images. We postulate that such a system processes barrier information. (Responding to data from Collett and Udin [1983], House has developed a different model to deal with worm information. This model builds on the earlier model of Didday of Section 5a to

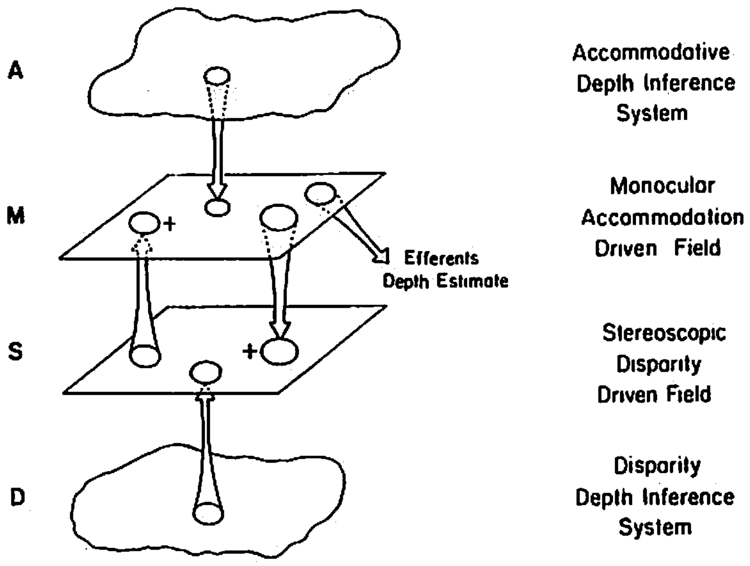

10. A depth mapping model using cross-coupling between an accommodation-driven system and a disparity-driven system. [House 1982]

combine prey-selection with prey-localization.) Figure 11f shows the outcome of such interaction. 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 behaviour of the animal.

While the above model is of interest in its own right — both as a model specific to the study of the amphibian, and as an indication of the class of stereopsis models based on multiple cues — we want to close this section by stressing 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 behaviour of the organism. Since we explicitly designed the model of Figure 10, we know that one layer represents accommodation information, while another represents disparity information — and we can clearly see the differences in these types of representation in Figure 11b. However, Figure 11f represents the sort of state of activity that is much more likely to be seen during the ongoing behaviour of the system, and here we see that both surfaces represent pooled information based on the interaction between the layers, rather than representing information directly supplied by sensory systems. *This clearly indicates the dangers of experimentation based on feature analysis without related high-level modelling.* As we can see, feature analysis of Figure 11f would simply show cells responsive to information available at a specific depth and visual direction. Only a far more subtle 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.

We may build on such studies to offer a high-level model which addresses behavioural data on depth and detour behaviour in toads in terms of interacting schemas to explain the spatial distribution of response, although it does not represent the details of a temporal processing. Figure 12 shows data of Collett (1982). Each dot-dash line pair indicates a position of the toad, with the dot indicating the position of its head, and the line indicating

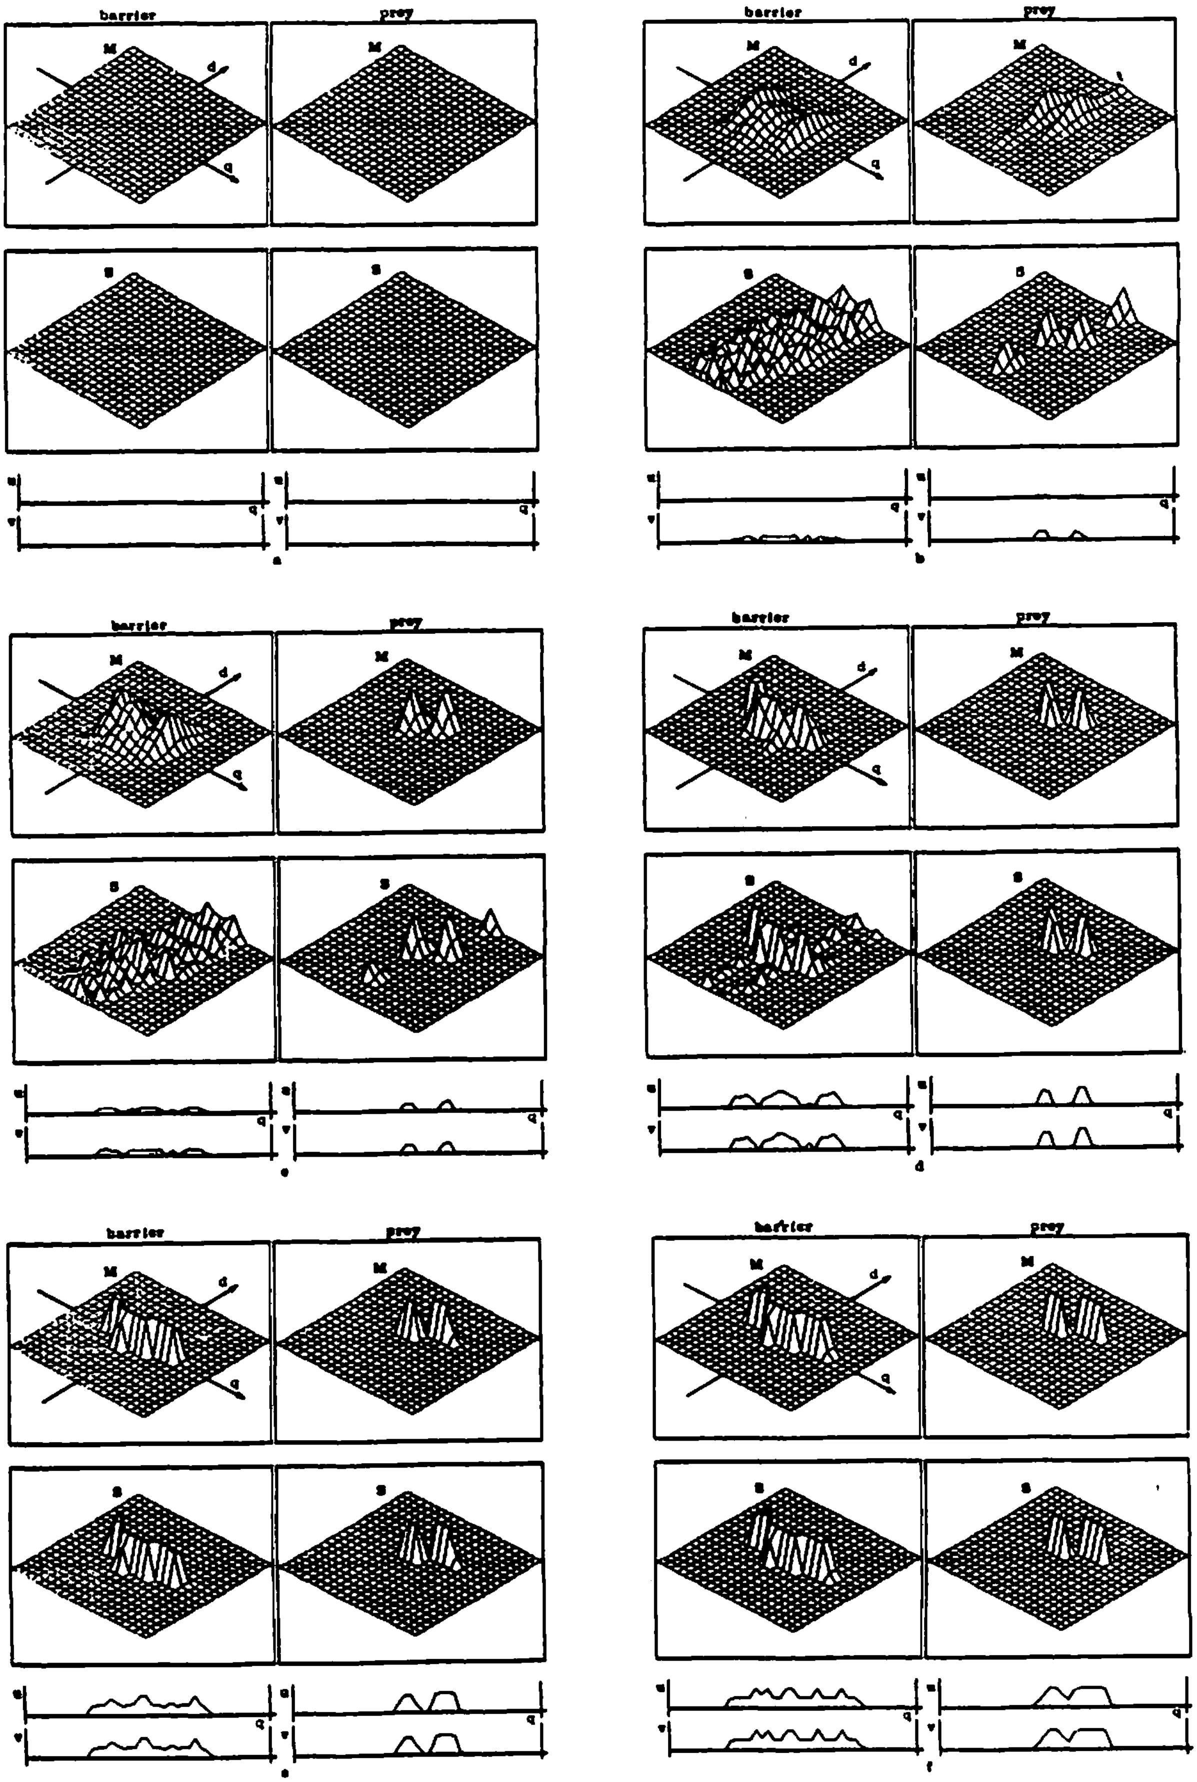
barrier
prey
M
S
d
q
u
v
a
b
c
d
e
f

11. Time Course of the Model — Base Case: 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 vs. disparity coordinate system. Successive figures are temporally spaced 1-1/3 field time-constants 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 and localization along the retinal angle axis of excitation in the inhibitory pools. [House 1982]

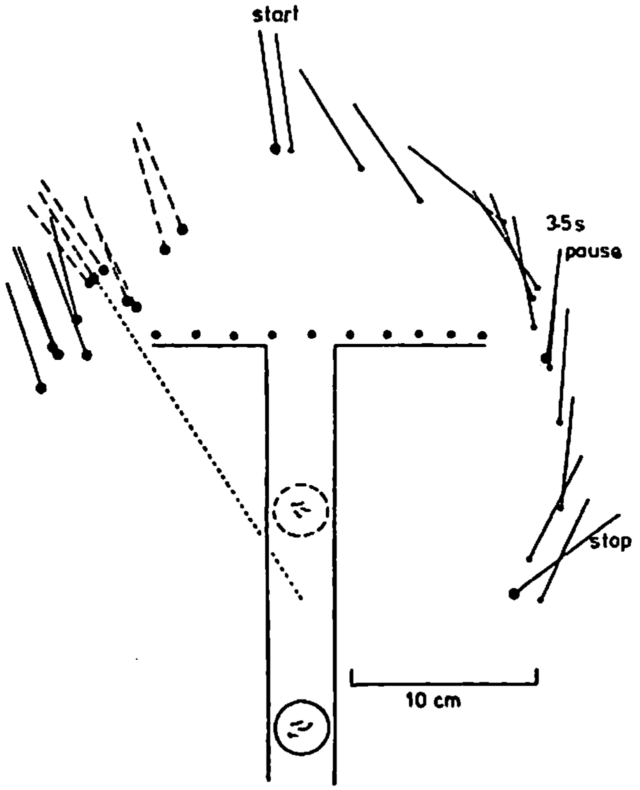

12. A trajectory of a toad which has sighted a number of mealworms behind a paling fence (the row of dots) and then detours around the fence to approach the prey. Note that when it stops (the prey no longer being in view) its position shows that it has retained a representation of prey position, which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

its orientation, as seen from above. The whole sequence represents one behaviour of the animal in response to the stimulus shown below it. The row of dots represents a paling fence, the T represents an opaque barrier, while the circle with squiggles in it represents a dish of mealworms. The mealworms are initially visible to the toad and lead, in combination with the fence, to the whole sequence of behaviour 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 the position of the worms. However, the final approach is aborted by the lack of adequate stimuli.

Given that the behaviour of the toad — whether to approach the prey directly, or to detour around the barrier — depends upon the distance at which the worms are behind the barrier, a full model of this behaviour must incorporate an analysis of the animal's perception of depth. An overall view of a possible "depth and detour" model is offered in Figure 13. Here we see that the retinal output of both eyes is separately processed for "barrier" and "worm" recognition, and that these are then combined to provide separate depth mappings for barrier and worm. We suggest that the animal's behaviour reflects the combined effects of prey "attraction" and barrier "repulsion". Formally, the barrier map B is convolved with a mask I which provides a (position-dependent) inhibitory effect for each fencepost; while the worm depth map W is convolved with a mask E which provides an excitatory effect for each worm. The resultant total map $T = B * I + W * E$ is then subject to further processing which will determine the chosen target, as well as providing other data which can be passed to coordinated motor schemas for side-stepping, orienting, jumping, and snapping, etc., to control the actual behaviour of the animal. Two alternative refinements of this overall schematic are offered by Arbib and House (1983), together with suggestions about how they might be subjected to experimental test. Lara et al. [1984] have suggested an alternative schema-based model. These alternatives provide an interesting challenge for future theory and experiment.

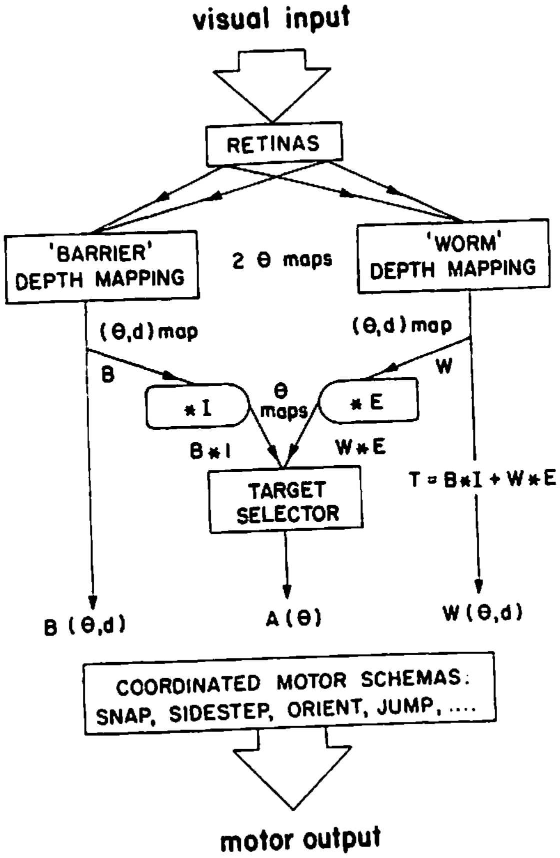

13. Conceptual Schematic of Visual/Motor Pathway: Assumptions made in this diagram are 1) that separate depth maps are maintained for prey and barrier stimuli, 2) direction for an orientation turn is obtained by combining information from these two depth maps, and 3) information on preferred orientation 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 and House 1985]

7. Tectal Column Models of Rana Computatrix.

Lettvin et al. found that group 2 retinal cells responded best to the movement of a small object within the receptive field; while group 3 cells responded best to the passage of a large object across the receptive field. It became common to speak of these cells as "bug detectors" (following Barlow [1953]) and "enemy detectors", respectively, though subsequent studies make it clear that the likelihood of a given frog behaviour will depend on far more than 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 the tectum as, inter alia, directing the snapping of the animal at small moving objects — it being known that the frog would ignore stationary objects, and would jump away from large moving objects. However, this notion of a simple stimulus-response chain via the tectum was vitiated by Ewert's observation that after a lesion to PT (pretectum-thalamus) a toad would snap at moving objects of all sizes, even those large enough to elicit escape responses in the normal animal. More detailed neurophysiological studies support the inference that the tectum alone will elicit a response to all (sufficiently) moving objects, and that it is PT-inhibition that blocks this response when the object is large, since tectal cells respond to visual presentation of large moving objects in the PT-lesioned animal [Ingle, 1973].

With this background, we may turn to data on toad behaviour for another perspective on top-down analysis of behavioural mechanisms in terms of interacting schemas. We may start with two basic behaviours of the animal — it will snap at small moving objects, and it will jump away from large moving 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 14) that animals with lesions of the pretectum will snap at large moving objects which a normal toad will avoid. This suggests a new analysis in terms of a prey-selection schema, which can be activated by moving objects of any size; and a predator-recognition schema, which serves not only to activate avoidance behaviour but also to inhibit prey acquisition behaviour.

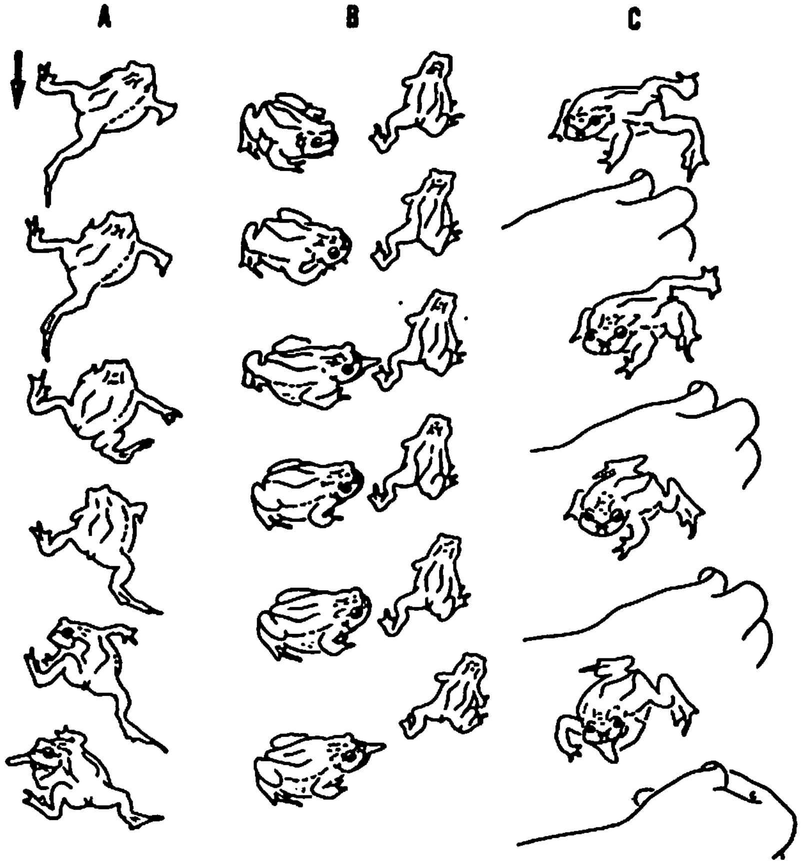

14. Three sequences showing a toad with pretectal lesion snapping at moving objects at which the normal toad would not snap ((a) its own foot; (b) the foot of another toad; (c) experimenter's hand). [Ewert 1976]

Thus, even gross lesion studies can distinguish between alternative top-down analyses of a given behaviour. Of course, such an analysis can be refined by more detailed behavioural studies which let us determine what features of a moving object serve to elicit one form of behaviour or another. For example, Ewert (Figure 15) placed a toad within a perspex cylinder from which it could observe 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 "prey-like" the object was) for different stimulus objects. As we can see from Figure 15(a), the worm-like stimulus (rectangle moved in the direction of its long axis) proved increasingly effective with increasing length; whereas even for 8 degrees of extension on its long axis, the antiworm stimulus (rectangle moved in the direction orthogonal to its long axis) proved ineffective in releasing orienting behaviour. The square showed an intermediate behaviour, the response it elicits rising to a maximum at 8 degrees, but being extinguished by 32 degrees. With such quantitative data to hand, Ewert and von Seelen (1974) produced the top-down model shown in Figure 16. As seen in Figure 16A, they postulated that retinal output was passed both to a tectal "worm filter", and a thalamic "antiworm filter", with the output of the latter serving to inhibit tectal (type II) activity which was excited by the former. As we can see in Figure 16B, a worm stimulus would then tend to yield much excitation of the worm filter which would be little inhibited by the thalamic antiworm response, thus yielding a vigorous output; while the antiworm (Figure 16D) would yield weak tectal type I activity, strong thalamic activity, and resultant weak tectal output. The square would yield intermediate behaviour. Ewert and von Seelen were able to adjust the parameters in this model to fit the data over a linear subrange of the results. However, our main criticism of the model is not that it is restricted to be linear, but that it is "lumped" in both space and time. That is, while 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 in the development of unlumped models which can indeed explain the spatio-temporal distribution of the animal's behaviour. To initiate this study, we first analyse the behaviour of a single tectal column.

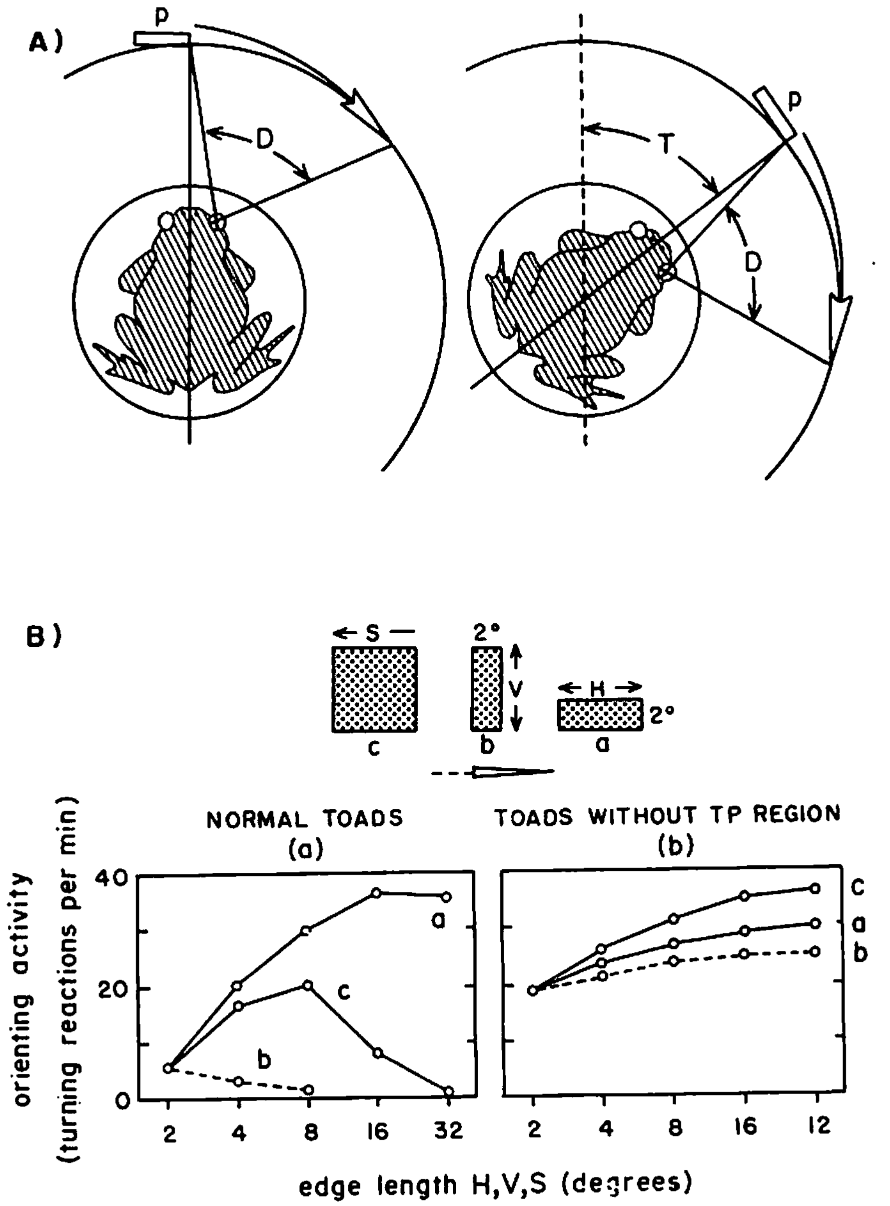

15. Prey-catching orienting behaviour to different configurations of the stimulus. A) Turning reaction to the stimulus presentation. D: effective angular displacement of the stimulus (p); T: angle of turning movements. B) Orienting activity to three stimulus configurations, horizontal ("worm": type a) and perpendicular ("anti-worm": 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].

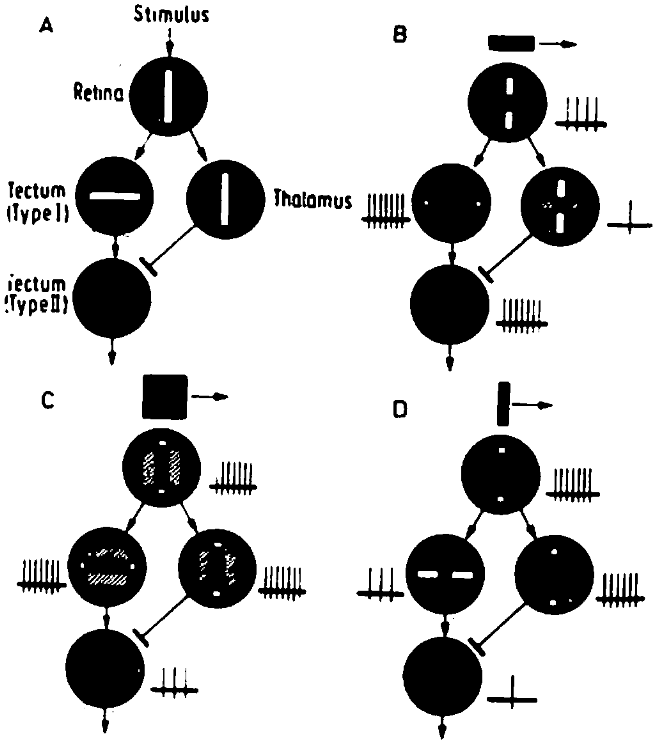

16. A lumped model of prey-predator discrimination: a worm filter provides excitatory input to an output cell which also receives inhibitory input from an antiworm cell. The result is that (B) a worm provides a potent response; (C) a square provides an intermediate response; while (D) an antiworm yields little or no response. [Ewert and Von Seelen 1974]

7a. *Facilitation of Prey-Catching behaviour.*

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 were presented initially for 0.3 sec, then removed, and then restored for only 0.3 sec, the second presentation would suffice to elicit a response, so long as the intervening delay was at most a few seconds. Ingle observed tectal cells whose time course of firing accorded well with this facilitation effect (Figure 17d). This leads us to a model [Lara, Arbib and Cromarty, 1982] in which the "short-term memory" is in terms of reverberatory neural activity rather than in terms of the short-term plastic changes in synaptic efficacy demonstrated, for example, by Kandel [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 known local neuroanatomy of the tectum. Unfortunately, the current knowledge of tectal circuitry is scanty, and much of the structure of the tectal column to be postulated below is hypothetical, and is in great need of confrontation with new and detailed anatomy and neurophysiology. The model described in this section addresses facilitation at a single locus of tectum. Further developments address the interaction of a number of columns, and we shall discuss these in Sections 7b and 7c.

As we discussed in Section 4, the tectal column model (Figure 7d) is abstracted somewhat crudely from the anatomy of Szekely and Lazar. It 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 17 and 18 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 7d are essentially the same.) All cells are modelled as excitatory, save for the stellates. The retinal input to the model is a lumped "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

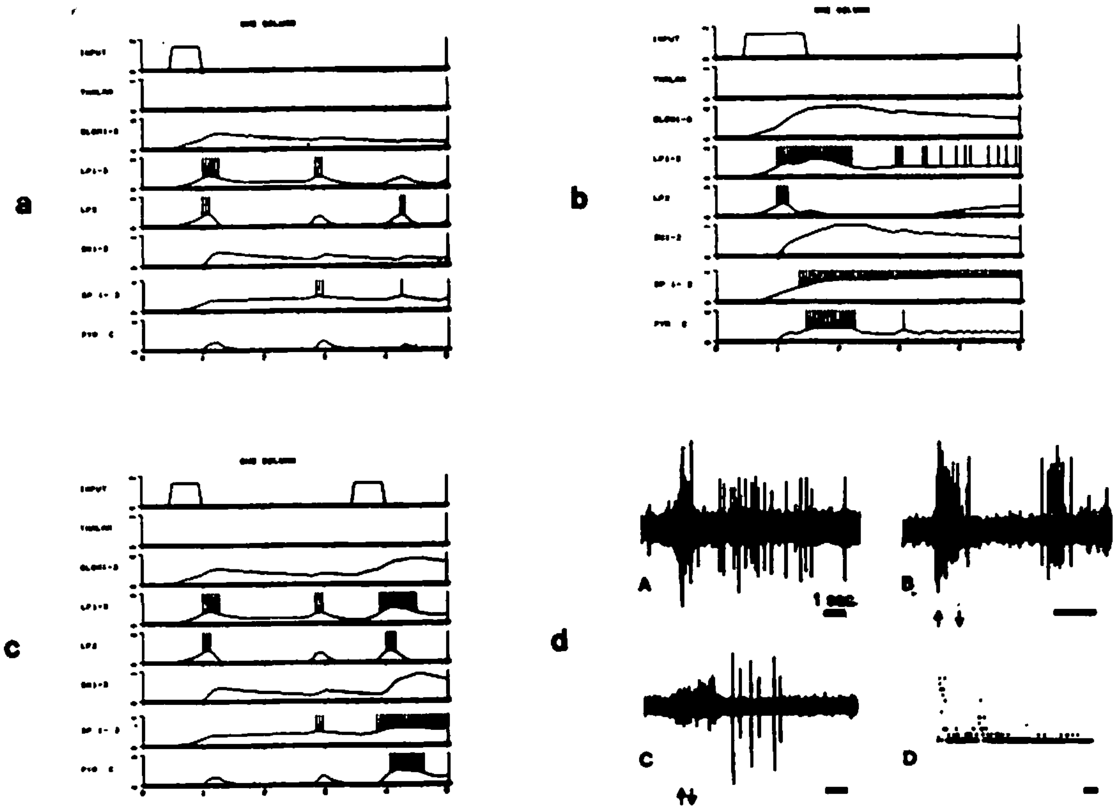

17. (a) Computer simulation of tectal cells response when a brief stimulus is presented. The onset of the stimulus produces a long lasting depolarizatin in the glomerulus which then fires the large pear-shaped cell (SP). This neuron in turn sends recurrent axons to the glomerulus and the stellate cell (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 frequency 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, Arbib and Cromarty 1982]

(d) Physiological behaviour of cells related to prey catching facilitation. A shows a brief class 2 burst followed by a delayed response of a tectal cell. In B the behaviour of a tectal cell is shown, responding to the presentation of the stimulus and again with a delay. C shows a tectal neuron that produces a delayed response to the presentation of the stimulus. Finally, D shows the postimulus histogram of a tectal cell showing a delayed peak at 3 to 4 seconds [From Ingle 1975].

inhibition. (For a full presentation of the differential equations used in the simulation, see Appendix 1 of Lara, Arbib and Cromarty [1982].)

The role of SN in our tectum model is reminiscent of Purkinje inhibition of the positive feedback between cerebellar nuclei and reticular nuclei, a basic component of Boylls' [1975, 1976] model of cerebellar modulation of motor schemas. Tsukahara [1972] found 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.

Returning to the tectal model: glomerular activity also excites the SP cell which also sends its axon back to the glomerulus. The SP cell also excites the LP cell to recruit the activity of the column. The PY cell is excited by both SP cell and LP cell. Clearly, the overall dynamics will depend upon the actual choice of excitatory and inhibitory weights and of membrane time constants. It required considerable computer experimentation to find the weights that yielded the neural patterns discussed below. Further study was devoted to a sensitivity analysis of how weighting patterns affect overall behaviour. It is our hope that our hypotheses on the ranges of the parameters involved in the model will stimulate more detailed anatomical and physiological studies of tectal activity.

Excitation of the input does not lead to runaway reverberation between the LP and its glomerulus; rather, this activity is "chopped" by stellate inhibition and we see a period of alternating LP and SN activity. The SP cell has a longer time constant, and is recruited only if this alternating activity continues long enough.

In one simulation experiment, we graphed the activity of the pyramidal cell as a function of the time for which a single stimulus is applied (Figure 18a). There is, as in the experimental data, a critical presentation length below which there is no pyramidal response. Input activity activates the LP, which re-excites 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 build-up of LP-SN alternating firing. If

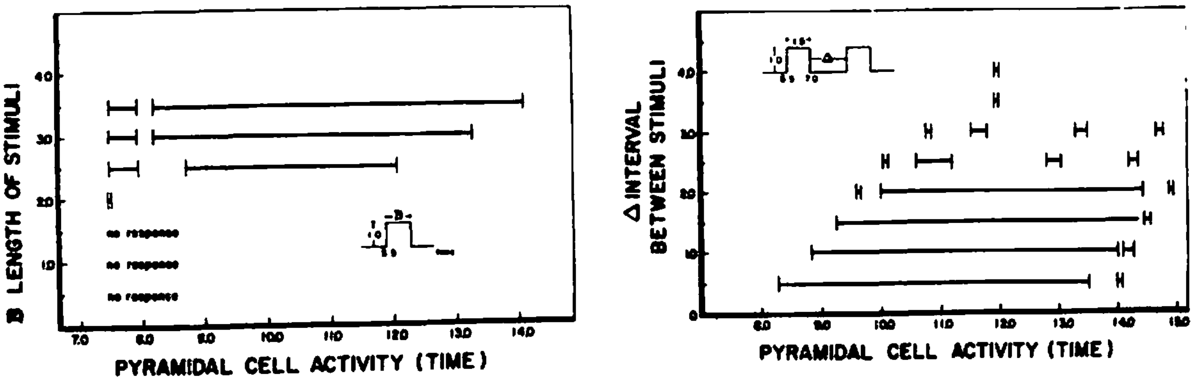

(a)

(b)

18. (a) Computer simulation of the PY behaviour when stimuli are presented for different intervals. (b) Computer simulation of the temporal pattern of the facilitation process after the presentation of a brief stimulus. [Lara, Arbib and Cromarty 1982]

the input is removed too soon, the reverberation will die out without activating the SP cells enough for their activity to combine with the LP activity and trigger the pyramidal output. However, if input is maintained long enough, the reverberation may continue, though not at a level sufficiently high to trigger output. However, a second simulation experiment (Figure 18b) shows that re-introduction of input within a short time after cessation of this "subthreshold" length of input presentation can indeed "ride upon" the residual reverberatory activity to build up to pyramidal input after a presentation time too short to yield output activity on an initial presentation.

7b *Worm-Antiworm Discrimination.*

With this, we now return to Ewert's behavioural data on prey-predator discrimination (Figure 15), and confront the task of explicating such data 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 15. In Figure 19a, we see that TH3 neurons (in thalamus/pretectum) have a response that is uniform for worms, increases with increasing length for antiworms, but is greatest for squares, increasing with their size. The activity of tectum type I neurons (T5(1)) seems to correlate fairly well with increase in 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 with the overall frequency of *behavioural* response; and it is this lumped neural response which is explained by the Ewert-von Seelen model of Figure 16. We note, too, that Figure 15b shows that the discrimination between worm, antiworm, and square is essentially abolished by pretectal lesion. Our task now is to give an unlumped model. As already discussed above, we want to do this in a way which is "evolutionary" in that — rather than seeking a minimal model of tectum to explain the specific pattern recognition task at hand — we seek to model the tectum in terms of the basic columnar structure already presented. The structure of the model is then that shown in Figure 20 (Cervantes-Peres, Lara and Arbib, 1985). Rather than using a single column to represent the tectum, we now use an 8x8 array of columns, and we provide the column not only with the R2 input of our initial study of facilitation,

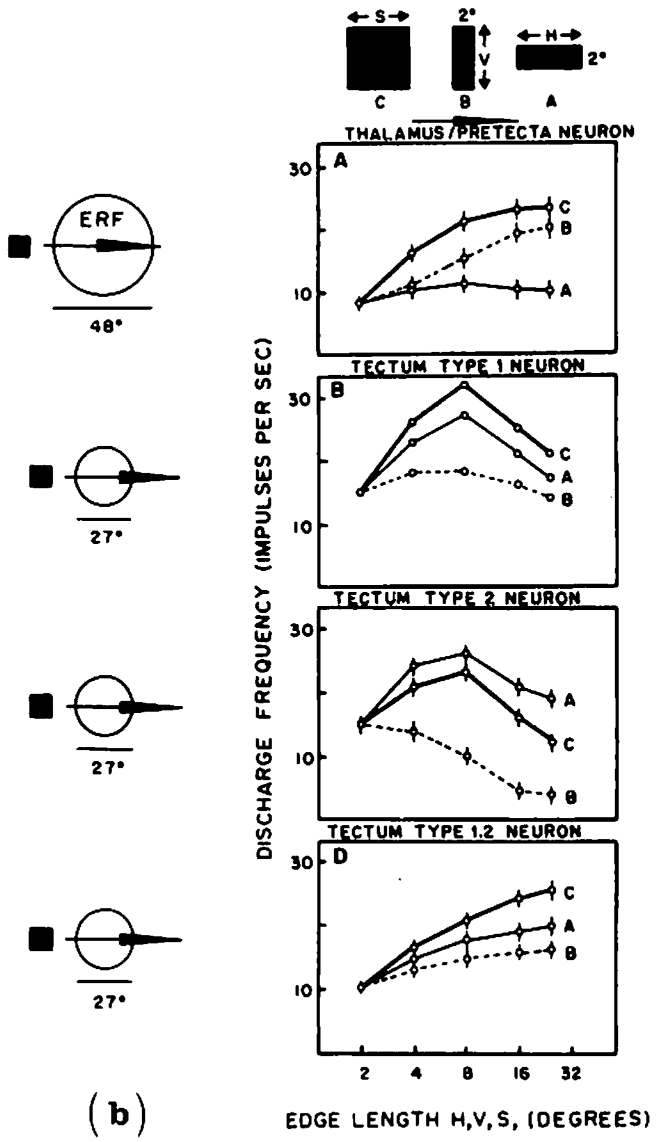

(b)

19. Tectal and Pretectal cell response from common toads to different configurations of moving stimuli (see Figure 8 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 prefers mostly stimuli type a, then type c, and gives a very weak response to type b. This neuron's response resembles the animal's behaviour (see Figure 2.B). 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].

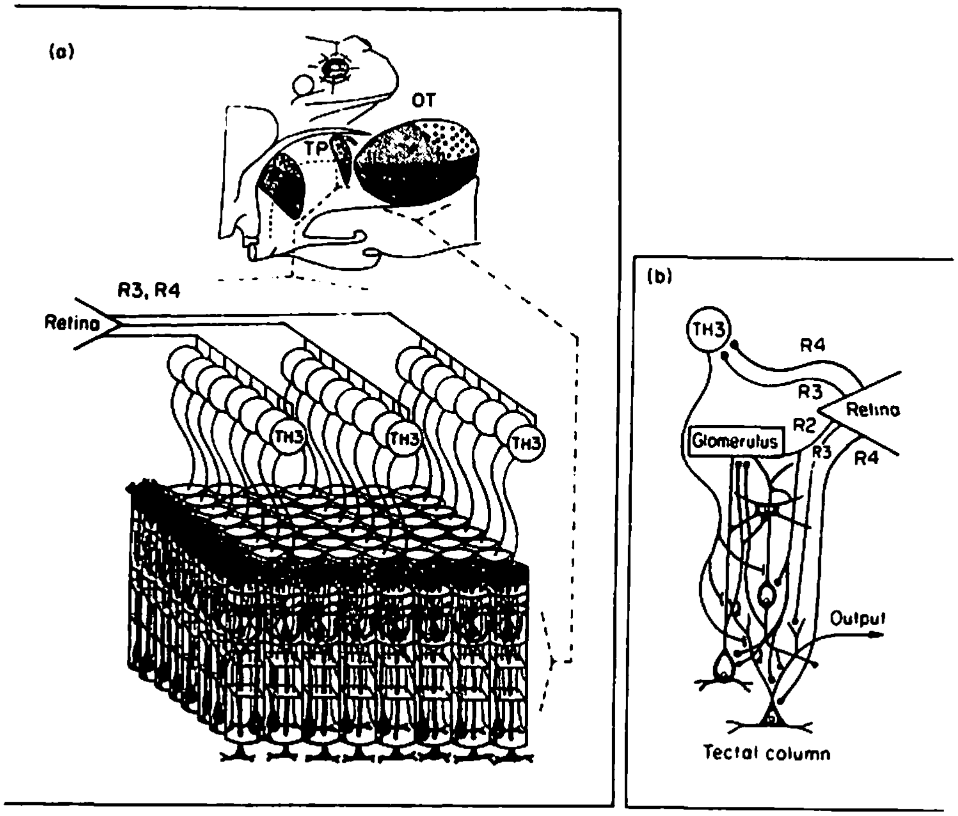

20. Interactions among retina, optic tectum and pretectum. The retina sends fibres in a retinotopical fashion to both optic tectum (class R2,R3 and R4), and pretectum (class R3 and R4). A) TH3 neurons also project retinotopicaly to the optic tectum. For simplicity we only show the projection of three rows of TH3 cells projecting upon the tectal columns. B) A closer look of 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, Lara and Arbib 1985].

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 response curves of Figure 21 which gave a lumped picture of ganglion cell response. (Current research is aimed as studying form-function relations in retina to derive a spatiotemporal pattern of retinal ganglion cell firing in response to varied activity.) The connections of R3 and R4 to TH3 are tuned to yield the responses shown in Figure 19a. With appropriate setting of parameters, the model does indeed exhibit, in computer simulation, responses to moving stimuli of different types (Figure 22) that match well the neural data of Figure 19. 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 general understanding of pattern recognition networks — further work must be done on tuning the model parameters. Figures 23 and 24 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 22. Figure 23 shows the response, for a brief period of time, of the 8x8 array of tectal columns to a moving worm (A), antiworm (B), and square (C) when no pretectal influence is included in the model. Each figure represents an 8x8 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 also yields more vigorous response on a cell-by-cell basis. Finally, Figure 24 shows the response of the full model in which the tectum is subject to pretectal inhibition. Here we see that there is 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 the paper by Cervantes-Perez, Lara, and Arbib (1985).

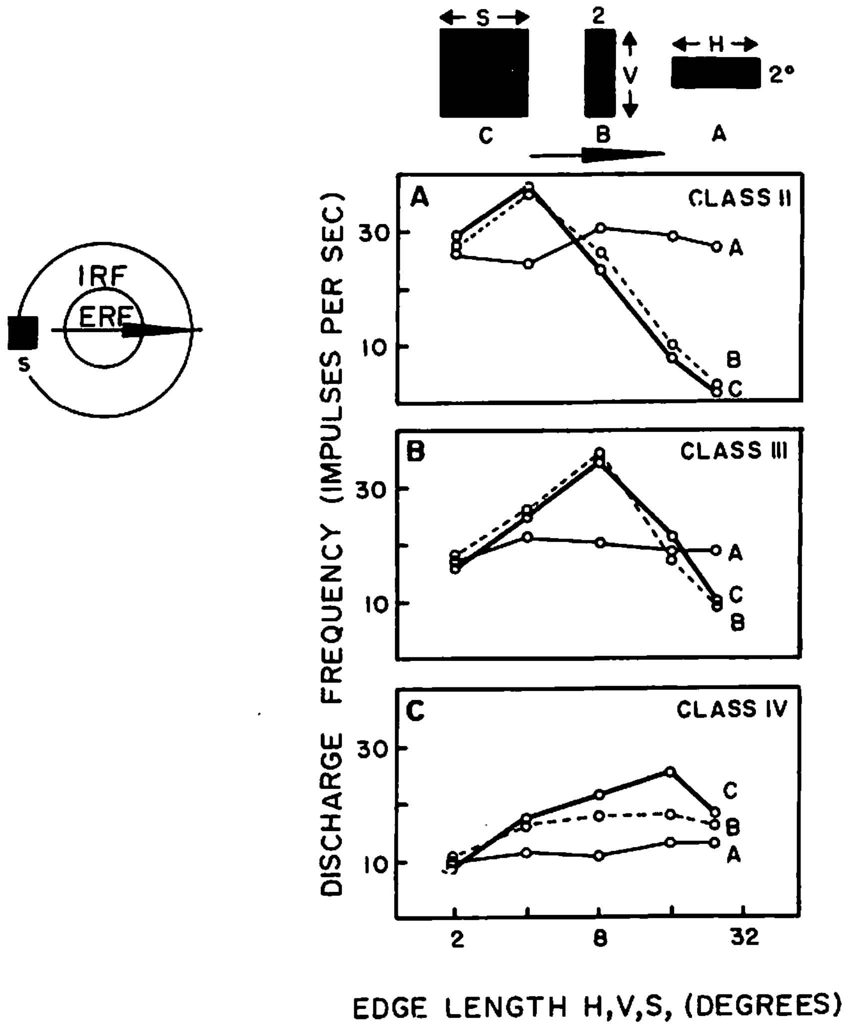

21. Retinal ganglion cells response (classes R2, R3, and R4) to different configurations (type a, b and c) of moving stimuli with a visual angular velocity of 7.6 degrees/sec. Left, they present a receptive field formed by a central excitatory (ERF) and a peripheral inhibitory (IRF) area. The three ganglion neurons respond almost with the same intensity to stimuli type a 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].

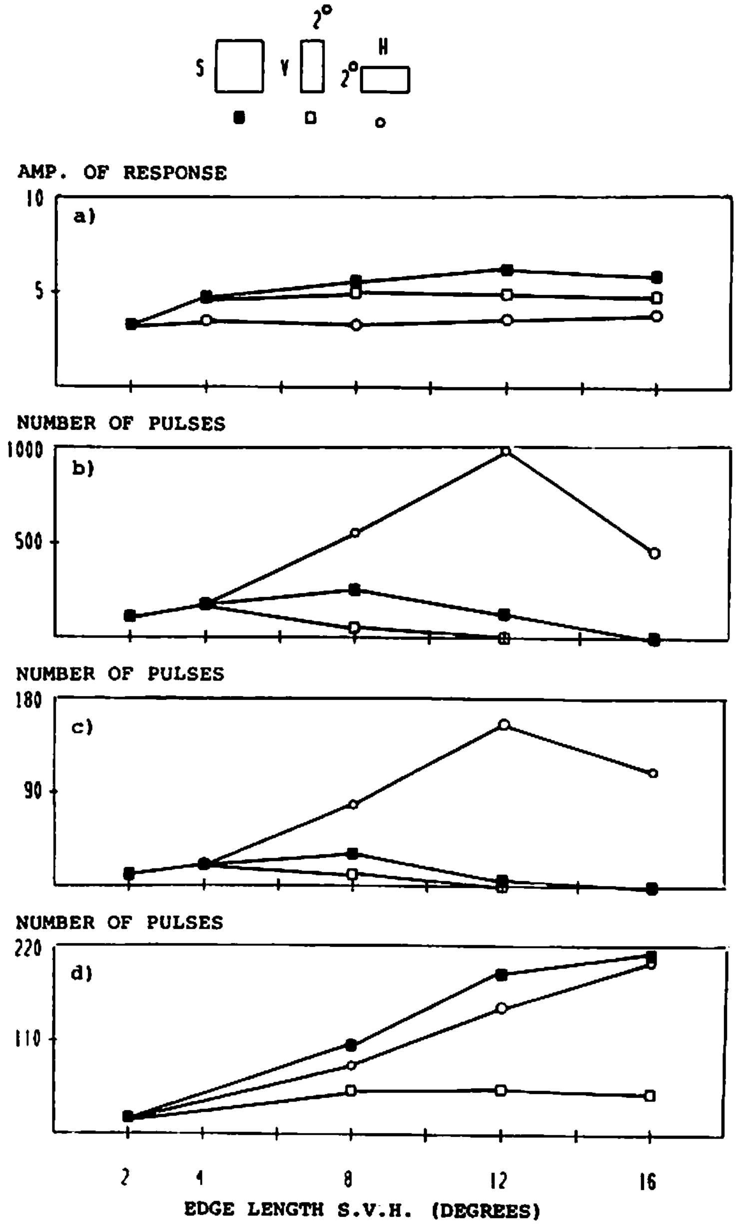

22. Computer simulation of the response of pretectal and tectal cells to different configurations (type a,b and c) of moving stimuli (visual angular velocity of 8 degrees/sec). A) Pretectal TH3 cell response: it is mostly sensitive to stimuli type c and b. B) Overall response of the tectum to the three types of stimulus (a, b and c): tectum response is mostly sensitive to stimulus type a and it is weaker to stimulus type b. C) Response of a PY cell to the three different stimuli : it respond 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 pretectum ablation occurs: these cells are mostly sensitive to stimuli of type c or a, and less to those of type b [From Cervantes-Perez, Lara and Arbib 1985].

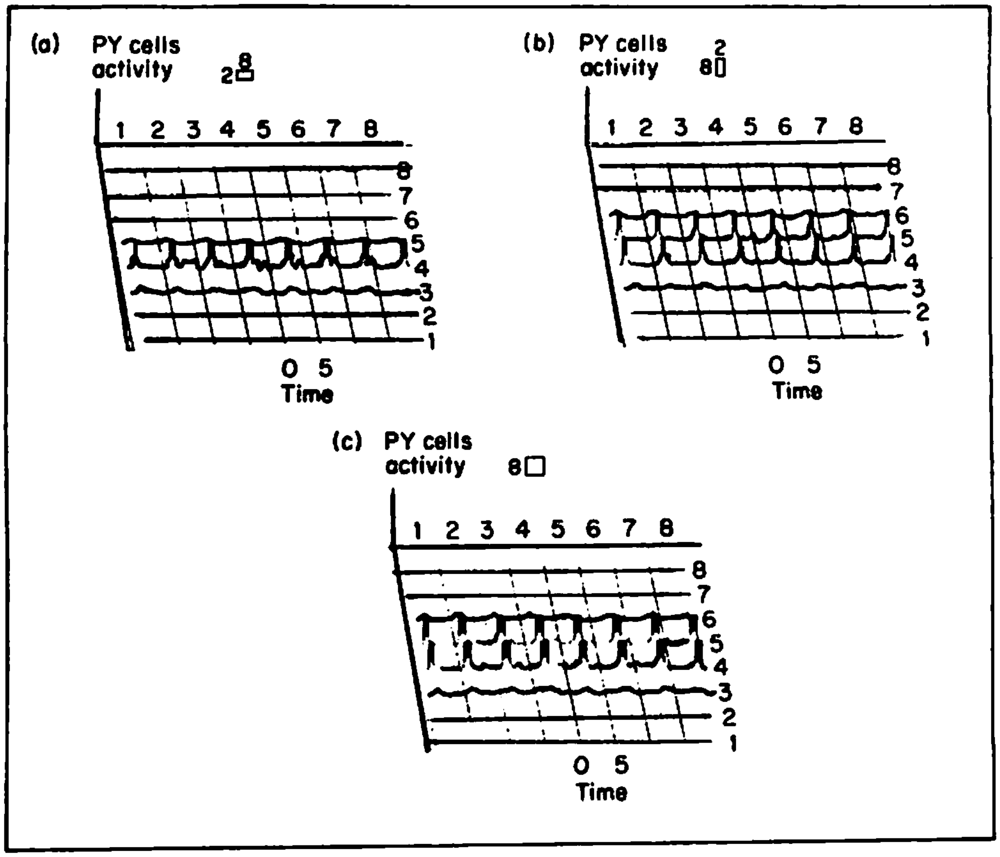

23. A computer simulation of the PY cell response of the 64 columns of the tectum to the different configurations of moving stimuli when pretectal ablation occurs: All stimuli are moved from left to right with a visual angular velocity of 8 degrees/sec. The figure shows four-dimensional graphs, where the x and y axes are used to represent the spatial localisation 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, while the vertical axis (z axis) represents its local membrane potential. In the graph when the PY local potential is about the threshold, this is indicated by spikes. A) response to a "worm-like" stimulus (type a) of 8x2 degrees; B) to the same stimulus moved as "antiworm-like" (type b); and C) to a square stimulus (type c) of 8x8 degrees. At the level of one PY cell, the tectal response is stronger for stimulus type c, then to type a, and finally to those of type b. The overall response of the tectum is also stronger to type c stimulus. It is also wider spread, while the response to "worm-like" stimulus, although it is strong too, is concentrated in a narrower area [From Cervantes-Perez, Lara and Arbib 1983].

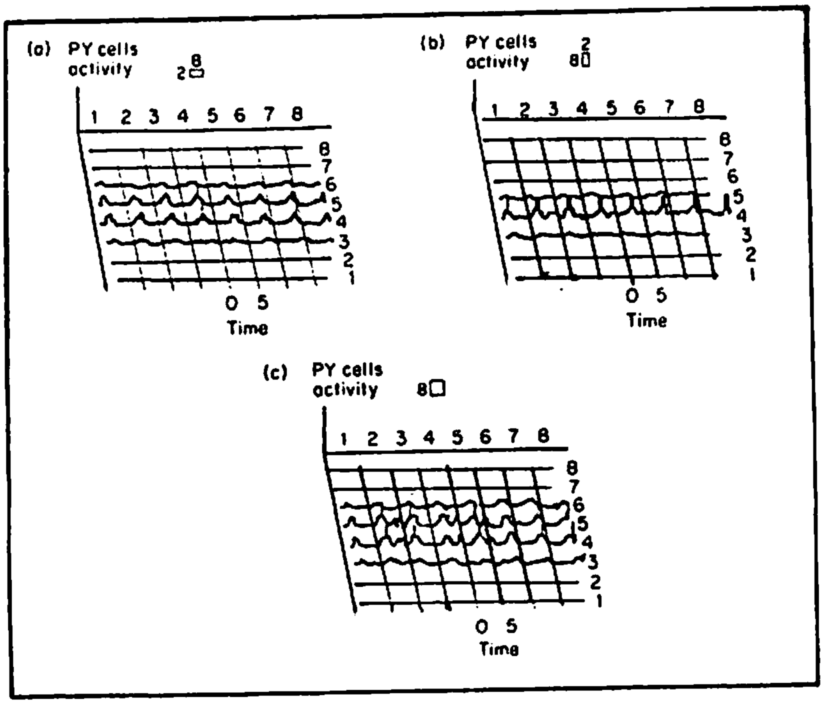

24. A computer simulation of the PY cell response of the 64 columns of the model of the interactions among retina, optic tectum and pretectum. All stimuli are moved from left to right with a speed of 8 degrees/sec. (See Figure 13 legend for an explanation of the graph characteristics). A) Response to a "worm-like" stimulus of 8x2 degrees (type a); B) Response to an "antiworm-like" stimulus of 2x8 degrees (type b); and C) Response to a square stimulus of 8x8 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 type c stimulus, although wider spread, is weaker than that to type a, whereas the weakest response is again to stimulus type b. PY cells are most sensitive to "worm-like" stimuli, rather than to squares or to "antworm-like" stimuli [From Cervantes-Peres, Lara and Arbib 1985].

Ewert (1976) has shown that, in frogs and toads, worm-antiworm discrimination is invariant to the velocity function of the stimulus, but Luthardt and Roth (1979) have reported that in *salamandra salamandra* the worm-antiworm discrimination varies with the velocity, the animal preferring "worm-like" to "antiworm-like" stimuli at low velocities, while at high velocities the "antiworm-like" stimulus is more effective. However, Himstedt (1982) argues that this phenomenon is not observed in all salamanders and that it probably depends on the animal's experience with certain types of prey, while Ewert has not found this change of preference. These observations lead us to two conclusions with the type of modelling presented in this section: first, the model is a generic model which must be specialized in different ways in the light of differing ethological, anatomical, and physiological observations on different species; and, second, the experimental literature at any time is not only incomplete, but is also not free from disagreements.

The job of the modeller, then, cannot be to provide a model constrained entirely by the current set of parameters, nor can or should modelling wait till "all the parameters are in". Rather our task is to closely work with experimentalists in providing models which better constrain the data, so that we can better assess which families of data can be integrated in a consistent fashion. With respect to the issues raised in this section, further developments of our model will require a more refined retina model as well as more detailed modelling of the temporal response of the different types of tectal neurons. Based on early simulations varying the membrane constant of the TH3 pretectal cell, we conclude that changes in an animal's experience that modify the pretectal inhibitory effect upon tectal activity could produce changes of the tectal preference between "wormlike" and square stimuli. Both experimental and simulation research addressing this hypothesis are necessary and will clarify the real nature of these processes.

References

Amari, S. (1982) Competitive and Cooperative Aspects in Dynamics of Neural Excitation and Self-Organization. In, *Competition and Cooperation in Neural Nets* (S. Amari and M.A. Arbib, Eds.), *Lecture Notes in Biomathematics*, Springer-Verlag.

Amari, S. and Arbib, M.A. (1977) Competition and cooperation in neural nets. In, *Systems Neuroscience* (Metzler, J., ed.), pp. 119-165, Academic Press, New York.

Apter, J.T. (1945) Projection of the retina on the superior colliculus of cats, *J. Neurophysiol.* 8: 123-134.

Apter, J.T. (1946) Eye movements following strychninization of the superior colliculus of cats, *J. Neurophysiol.* 9: 73-85.

Arbib, M.A. (1975) Artificial intelligence and brain theory: unities and diversities, *Ann. Biomed. Eng.* 3: 238-274.

Arbib, M.A. (1981) Perceptual structures and distributed motor control. In, *Handbook of Physiology-The Nervous System II. Motor Control* (ed. V.B. Brooks), Bethesda, MD: Amer. Physiological Society. pp. 1449-1480.

Arbib, M.A. (1982a) *Rana Computatrix*: An evolving model of visuomotor coordination in frog and toad. In: *Machine Intelligence* 10 (J.E. Hayes, D. Michie and Y. H. Pao, Eds.), Chichester: Ellis Horwood, pp. 501-517.

Arbib, M.A. (1982b) Modelling neural mechanisms of visuomotor coordination in frog and toad. In, *Competition and Cooperation in Neural Nets* (S. Amari and M.A. Arbib, Eds.), Lecture Notes in Biomathematics Vol. 45, Springer-Verlag, pp. 342-370.

Arbib, M.A., Boylls, C.C., and Dev., P., (1974) Neural Models of Spatial Perception and the Control of Movement. In, *Cybernetics and Bionics*, (W.D. Keidel, W.Handler, and M. Spreng, Eds.), Munich: Oldenbourg, 216-231.

Arbib, M.A. and House, D. (1983) Depth and detours: Towards neural models, *Proceedings of second workshop on visuomotor coordination in frog and toad: Models and Experiments* (R. Lara and Arbib, M.A., Eds.) COINS-Technical report 83-19, University of Massachusetts, Amherst, Massachusetts.

Arbib, M.A., Iberall, T. and Lyons, D. (1985) Coordinated control programs for movements of the hand, *Exp. Brain Res. Suppl.* 10, pp. 111-129.

Barlow, H. (1953) Summation and inhibition in the frog's retina, *J. Physiol. (Lond.)* 119: 69-88.

Barlow, H.B., Blakemore, C., and Pettigrew, J.D. (1967) The neural mechanism of binocular depth discrimination, *J. Physiol.* 193: 327-342.

Bartlett, F.C. (1932) *Remembering*. London: Cambridge University Press.

Bobrow, D.G., and A. Collins (1975) *Representation and Understanding: Studies in Cognitive Science.* New York: Academic.

Borges, J. L. (1975) Of exactitude in science. In, *A Universal History of Infamy*, p. 131, Penguin Books.

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.

Boylls, C.C. (1976) A Theory of Cerebellar Function with Applications to Locomotion. II. The relation of anterior lobe climbing fiber function to locomotor behaviour in the cat. Technical Report 76-1, Computer and Information Science Department, University of Massachusetts at Amherst.

Cervantes-Perez, F., Lara, R. and Arbib, M.A. (1985) A neural model of interactions subserving prey-predator discrimination and size preference in anuran amphibia, *J. Theoretical Biology* 113, 117-152.

Collett, T. (1977) Stereopsis in toads, *Nature* 267: 349-351.

Collett, T. (1982) Do toads plan routes? A study of the detour behaviour of *Bufo Viridis*, *J. Comp. Physiol.* 146: 261-271.

Collett, T., and 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* (R. Lara and Arbib, M.A., Eds). COINS-Technical Report 83-19, University of Massachusetts, Amherst, Massachusetts.

Craik, K.J.W.: (1943) *The Nature of Explanation.* London: Cambridge Univ. Press.

Dev, P. (1975) Perception of depth surfaces in random-dot stereograms: A neural model, *Int. J. Man-Machine Studies* 7: 511-528.

Didday, R.L. (1970) The Simulation and Modelling of Distributed Information Processing in the Frog Visual System, Ph.D. Dissertation, Stanford University.

Didday, R. (1976) A Model of visuomotor mechanisms in the frog optic tectum, *Mathematical Biosciences* 30:169-180.

Ewert, J.P. (1976) The visual system of the toad: behavioural and physiological studies on a pattern recognition system. In, *The Amphibian Visual System* (Fite, K., ed.), pp. 142-202. Academic Press, New York.

Ewert, J.P. and Von Seelen, W. (1974) Neurobiologie und System-Theorie eines Visuellen Muster-Erkennungsmechanismus bei Krote, *Kybernetik* 14:167-183.

Fredericks, J.A.M. (1969) Disorders of the body schema. In, *Handbook of Clinical Neurology. Disorders of Speech Perception and Symbolic Behaviour.* edited by P.J. Vinken and G.W. Bruyn. Amsterdam: North Holland, vol. 4, p. 207-240.

Gregory, R.L. (1969) On how so little information controls so much behaviour. In, *Towards a Theoretical Biology, 2: Sketches*, edited by C.H. Waddington. Edinburgh: Edinburgh University Press.

Head, H., and G. Holmes (1911) Sensory disturbances from cerebral lesions, *Brain* 34: 102-254.

Himstedt, W. (1982) Prey selection in salamanders In:, *Analysis of visual behaviour.* (D.J. Ingle, M.A. Goodale and R.J.W. Mansfield, Eds.), The MIT Press, Cambridge, pp. 47-66.

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

House, D. (1984) A Model of Prey-Localisation in Frog and Toad, Chapter IV, Ph.D. Dissertation, Department of Computer and Information Science, University of Massachusetts at Amherst.

Ingle, D. (1968) Visual releasers of prey catching behaviour in frogs and toads, *Brain, Behav., Evol. 1*: 500-518.

Ingle, D. (1973) Disinhibition of tectal neurons by pretectal lesions in the frog, *Science* 180: 422-424.

Ingle, D. (1975) Focal attention in the frog: behavioural and physiological correlates, *Science* 188: 1033-1035.

Ingle, D. (1976) Spatial visions in anurans. In, *The Amphibian Visual System* (K. Fite, Ed.), Academic Press: New York, pp. 119-140.

Jeannerod, M. (1981) Le controle de l'oeil sur le geste, *La Recherche*, No. 120, Vol. 12, March.

Jeannerod, M., Michel, F., and Prablanc, C. (1984) The control of hand movements in a case of hemianaesthesia following a parietal lesion, *Brain*, 107, 899-920.

Julesz, B. (1971) *Foundations of Cyclopean Perception* Chicago: Univ. of Chicago Press.

Kandel, E.R. (1978) A Cell-Biological Approach to Learning, Bethesda, MD, *Soc. Neurosci.*.

Lara, R., Arbib, M.A. and Cromarty, A.S. (1982) The role of the tectal column in facilitation of amphibian prey-catching behaviour: a neural model, *J. Neuroscience*. 2: 521-530.

Lara, R., Carmona, M. and Crus, A. (1984) A global Model of the Neural Mechanisms Responsible for Visuomotor Coordination in Toads, *J. Th. Biol 110,* 587-618.

Lettvin, J. Y., Maturana, H., McCulloch, W. S. and Pitts, W. H. (1959) What the frog's eye tells the frog brain, *Proc. IRE.* 47: 1940-1951.

Luthardt, G. and Roth, G. (1979) The relationship between stimulus orientation and stimulus movement pattern in the prey catching behaviour of *Salamandra Salamandra*, *Copeia* 3: 442-447.

MacKay, D.M. (1966) Cerebral organization and the conscious control of action. In, *Brain and Conscious Experience*, edited by J.C. Eccles. New York: Springer-Verlag, p. 422-440.

Marr, D., and Poggio, T. (1977) Cooperative computation of stereo disparity, *Science*. 194: 283-287.

Minsky, M.L. (1961) Steps towards artificial intelligence. *Proc. IRE* 49: 8-30.

Minsky, M.L. (1975) A framework for representing knowledge. In, *The Psychology of Computer Vision*, edited by P.H. Winston. New York: McGraw-Hill, p. 211-277.

Neisser, U. (1976) *Cognition and Reality: Principles and Implications of Cognitive Psychology.* San Francisco: Freeman.

Nelson, J. I. (1975) Globality and stereoscopic fusion in binocular vision, *J. Theor. Biol.* 49: 1-88.

Oldfield, R.C., and Zangwill, O.L. (1943) Head's concept of the body schema and its application in contemporary British psychology, *Br. J. Psychol.* 32: 267-286, 1942; 33: 58-6 , 113-129, 143-149.

Pettigrew, J.D., Nikara, T., and Bishop, P.O. (1968) Binocular interaction on single units in cat striate cortex, *Exp. Brain Res.* 6: 391-410.

Piaget, J. (1971) *J. Biology and Knowledge.* Edinburgh: Edinburgh Univ. Press.

Pitts, W.H., and McCulloch, W.S. (1947) How we know universals, the perception of auditory and visual forms, *Bull. Math. Biophys.* 9: 127-147.

Schmidt, R.A. (1975) A schema theory of discrete motor skill learning, *Psychol. Rev.* 82: 225-260.

Schmidt, R.A. (1976) The schema as a solution to some persistent problems in motor learning theory. In, *Motor Control: Issues and Trends*, edited by G.E. Stelmach. New York: Academic, p. 41-65.

Sperling, G. (1970) Binocular vision: a physical and a neural theory, *Am. J. Psych.* 83: 461-534.

Szekely, G. and Lasar, G. (1976) Cellular and synaptic architecture of the optic tectum. In, *Frog Neurobiology* (Llinas, R. and Precht, W., eds.). Springer Verlag, pp. 407-434.

Szentagothai, J. and Arbib, M. A. (1974) *Conceptual Models of Neural Organization. NRP Bulletin* vol. 12, no. 3: 310-479. (Also: The MIT Press, 1975).

Tsukahara, N. (1972) The properties of the cerebello-pontine reverberating circuit, *Brain Res.* 40: 67-71.