To appear in: SELFORGANIZING SYSTEMS, The Emergence of Order.
(F.E. Yates, Editor) Plenum Press, 1982.

A VIEW OF BRAIN THEORY*

Michael A. Arbib
Computer & Information Science Department,
and Center for Systems Neuroscience

COINS Technical Report 81-31

(Reprinted March 1982)
(and again October 1982)
(and a third time in March 1983)

---

* Portions of this paper were presented in two talks: "Perceptual Structures and Distributed Motor Control" delivered at the School of Social Sciences, University of California at Irvine, on January 29, 1980; and "The Theoretical and Experimental Bases of Brain Theory" delivered in the Department of Biomedical Engineering at the University of Southern California on February 14, 1980. I thank Ken Wexler and his colleagues and George Moore for their hospitality on these occasions.

Preparation of the paper was supported in part by NIH grant NS14971-02.

TABLE OF CONTENTS

Section 1 of the chapter distinguishes top-down brain theory (functional analysis of cognitive processes) from bottom-up brain theory (as in the analysis of the dynamics of neural nets). Section 2 then advocates 'cooperative computation' as providing a style of analysis for studying the interactions of neural subsystems.

Section 3, "Interacting Schemas for Motor Control," provides a top-down analysis of perception and the control of movement in the 'action-perception cycle'. We introduce perceptual schemas as the building blocks for the representation of the perceived environment, and motor schemas as control systems to be coordinated into programs for the control of movement. Section 4 contrasts two approaches to the design of machine vision systems. The analysis exhibits many of the insights to be gained from a top-down analysis, but shows that such an analysis does not guarantee a unique functional analysis of the problem at hand. Section 5 presents an algorithm for computing the optic flow which exhibits the style of cooperative computation set forth in Section 2. This model has not been confronted with data from neurophysiology, but is very much 'in the style of the brain,' and offers interesting insights into the evolution of layered neural structures.

The next two sections do offer neural models which have developed through a rich interaction between theory and experiment. Section 6 models the possible role of the cerebellum in parametric tuning of motor schemas; while Section 7 models interaction between tectum and pretectum in visuomotor coordination in frog and toad. We close in Section 8 by sketching a path between these neural models and the top-down analysis of cognition.

1. BRAIN THEORY: 'BOTTOM-UP' AND 'TOP-DOWN'

'Brain theory' seeks to make systematic use of mathematical analysis and computer simulation to elucidate the interactions of the components of the brain, and how they can subserve such diverse functions as perception, memory, and the control of movement. As such it differs from the 'mind theory' of cognitive psychology which seeks to analyze properties of the mind with little concern for how these properties are played out over the structures of the brain. In this paper, we provide a view of brain theory informed by two viewpoints: the need for a healthy interaction between cognitive studies ('top-down') and neuroscience ('bottom-up); and the emerging utility of an approach to brain theory which emphasizes the 'cooperative computation' of a multitude of subsystems. (The latter point is also central to the paper by von der Malsburg in this volume. We shall develop it in Section 2.)

I want to suggest that brain theory should confront the 'bottom-up' analyses of neural modelling not only with biological control theory but also with the 'top-down' analyses of artificial intelligence and cognitive psychology (Arbib, 1975; 1978; Marr and Poggio, 1977a). In bottom-up analyses, we take components of known function, and explore ways of putting them together to synthesize more and more complex systems. In top-down analyses, we start from some complex functional behavior that interests us, and try to determine what are natural subsystems into which we can decompose a system that performs in the specified way. I would argue that progress in brain theory will depend on the cyclic interaction of these two methodologies. In advocating a brain theory of this type, we suggest that many experiments in the laboratory of the neuroanatomist and the neurophysiologist should be related to evolving theories of high-level brain function, while at the same time suggesting that to contribute to brain research a cognitive analysis must constrain the subsystems posited by exploring whether they can be mapped into the circuitry of actual brain regions.

The top-down approach complements bottom-up studies, for one cannot simply wait until one knows what all the neurons are and how they are connected to then simulate the complete system. Jose Luis Borges (1975) tells of a country which prided itself on the excellence of its cartography. As years went by, the cartographers produced maps of greater and greater accuracy, until finally they achieved the ultimate, full-scale, map. And Borges wryly notes that there are places in the Western Deserts where even today you can see tattered fragments of the map (presumably, pegged to the place they represent). We need a guide to understand a new territory, but a map which provides no simplifications and pointers to distinctive features to aid our exploration does not help us. In the same way, a model which simply duplicates the brain is no more illuminating than the brain itself. We need theory to process data efficiently, and to present the facts in an illuminating and insightful way; and we need detailed studies by neurologists of how brain lesions impair behavior (both transiently and permanently), as well as by neurophysiologists and neuroanatomists on circuit, cell and synapse. Brain theory, properly conceived, can contribute to the design of further experiments that help shape the overall understanding of how portions of the brain interact to make us what we are.

The layered structure of the brain is one of its distinctive features, and may well play an important role in helping us analyze the way in which sensory information -- be it visual or somatic or otherwise -- can be used in the control of behavior. Such an analysis offers hope that rather than having to model individual activity in millions or more of neurons, we can hope to understand much brain function in terms of the interaction of spatial patterns distributed across a relatively small number of neural arrays. Another unit of complexity intermediate between the single neuron and the brain is the 'module'. One of the earliest 'modules' came from the Scheibels' (1958) anatomy of the reticular formation. They observed that the major neurons of the reticular formation had dendrites which were parallel to one another, and orthogonal to the axons which ran up and down along the head-tail axis. They thus suggested that nearby neurons could be aggregated together into 'poker chips' orthogonal to the head-tail axis, with the neurons within a module being roughly uniform in their sampling of the traffic up and down the reticular formation as well as their sampling of the peripheral input.

This analysis of the reticular formation in terms of the interactions between a relatively small number of modules was used by Kilmer and McCulloch (1969) in their RETIC model of the reticular formation. This model was interesting not only in that it was related to one of the earliest 'module' concepts within neuroanatomy, but also because it showed how a neural system could achieve some overall behavior without executive control. Kilmer and McCulloch suggested that the reticular formation had the task of committing the organism to some overall mode of behavior. Each individual module of RETIC was thus to use its sample of inputs to make an initial determination of the relative desirability of the different modes. Different modules were coupled to a sample of their neighbors in such a way that the back and forth traffic between the modules proved sufficient for them to reach a consensus in which the majority gave top priority to a single mode thus committing the organism overall. We shall say more about this style of 'cooperative computation' in Section 2.

Another form of intermediate unit is the 'column', an anatomical structure suggested first by the studies of Mountcastle (1957) and Powell on somatosensory cortex (1957), and then later by the work of Hubel and Wiesel on the visual cortex (1974). The suggestion here is that as we move up and down through the layers of the cortex we find that all the neurons are responsive to roughly the same stimuli from the external world, and that a small displacement will keep us with neurons with roughly the same 'features', but that if we go further we will move on to another 'column' of cells describable by different features. This suggests, then, that much of the analysis of cortex can be conducted in terms of the interaction between columns, with the analysis of individual neurons playing a more restricted role to explain the dynamics of the column units (Mountcastle, 1978; Szentagothai and Arbib, 1975; Szentagothai, 1978).

Before closing this section, we stress that brain theory will progress both by computer simulation and mathematical analysis. We see analyses that are mathematical in that we not only provide formal descriptions of systems but also prove theorems about their behavior. There are mathematical analyses of properties of general classes of systems related to the Hodgkin-Huxley equation, (reviewed in (Rinzel, 1978)), in studies of cooperative computation (as outlined in the next subsection), and in control theory. On the other hand, there are many cases in which our symbolic representation of systems and their interactions does not take a form which lends itself easily to mathematical analysis. Rather, we turn to the computer to conduct neural simulations and cognitive modelling experiments in the style of Artificial Intelligence (AI) to gain insight into the capabilities of a system so represented. It may well be that we shall see the growth of a delicate interaction between mathematical analysis and simulation as we use our mathematics to determine what is a sufficiently wide sample of different conditions in which to simulate a system to get a proper appreciation of the full range of its behavior.

2. COOPERATIVE COMPUTATION

Time and again, we find that modern studies in brain theory must concern themselves with the integration of the activity of a multitude of subsystems. This brings us, then, to the key question of cooperative computation: How is it that local interaction of a number of systems can be integrated to yield some overall result? The study of cooperative phenomena has its roots in the statistical mechanics of physics whereby the individual motion of billions of atoms could be averaged to come up with reliable thermodynamic descriptions of the behavior of a gas or a liquid as a whole. For example, in studying ferromagnets, we seek to understand how atomic magnets can 'cooperate' to yield global magnetism through the mass effects of local interactions. Cragg and Temperley (1953) were perhaps the first to suggest analogies between cortical activity and domain formation in ferromagnets. A number of interesting models have been developed in the 1970's, starting with the studies of Harth (1970) and of Wilson and Cowan (1973) and continuing up to such studies as that of Amari and Arbib (1977), Amari (1980), and Shaw (1978). For example, if we imagine the visual cortex to contain a great variety of cells tagged not only for visual direction but also for depth in the visual field, then we can imagine the process of recognizing regions in the visual input at different depths to be one of suppressing all neural activity except that corresponding to the depths within a given direction; and we can then imagine this process of segmentation as having much in common with the process of domain formation in magnets (Julesz, 1971). We shall say more about stereopsis below.

The study of how local interactions may yield global function may be studied 'discretely' or 'continuously'. The classic 'discrete' study is given by von Neumann's (1951) use of tesselation automata to model self-reproduction. These structures comprise nets of regularly spaced automata, each connected to a few neighbors, with the next state of each unit being determined by the present state of units in its neighborhood. (This will be familiar to readers of <u>Scientific American</u> as the setting for Conway's "Game of Life".) The analysis of such nets is purely combinatorial: given a program for local cellular interaction, we check that the details of the program do indeed yield some desired pattern formation. (See Arbib (1972b) and Ede (1978) for the use of such models in embryology.) By contrast, the 'continuous' approach approximates a tissue of cells by functions varying continuously over the tissue, and uses techniques from differential equations, stability, and statistics. This approach to biological systems goes back to Turing's (1952) paper on morphogenesis (which was fore-shadowed by Rashevsky (1948)), which introduced the use of reaction-diffusion equations into the study of pattern formation.

Turing studied a ring of cells. In each cell there were two chemical substances called morphogens. Within any one cell, these substances could engage in chemical reactions; each morphogen could also diffuse between cells. One might think that such interactions would yield identical chemical equilibria in all the cells. To the contrary, Turing was able to show that, even with linear equations, the system would eventually be structured with standing waves of chemical concentrations, thus providing the substrate for the expression of biological pattern. Turing was once asked whether his model would explain the stripes of the zebra, and his reply was "The stripes are easy, it's the horse part that I have trouble with!" The important point for us here is that we have a model of how local interactions give rise to global pattern (See (Fife, 1979; Katchalsky et al., 1974; and Kopell, 1978) for further information on reaction-diffusion equations and related topics.) Grossberg (1978) has explored analogies between reaction-diffusion equations and neural processes, including the problem of patterning an array of synapses in neural learning. Haken (1978) has written a book on "Synergetics," seeking to provide a unified mathematical framework in which a number of these cooperative phenomena can be viewed. In particular, he relates the reaction-diffusion problem to the mathematics used to look at such phase transitions as the formation of a coherent pulse of light in a laser and the order-disorder transition in a magnet.

In 'top-down' brain theory and artificial intelligence, we find a 'discrete' style of analysis of cooperative computation. Such AI projects as HEARSAY (Erman and Lesser [1980]) and VISIONS (Hanson and Riseman [1978a]) use a number of interacting knowledge sources to converge upon a perceptual analysis of some sensory input -- an acoustic signal encoding a sentence in the first case, and a color photograph of an outdoor scene in the second. The conceptual structure thus created seems to hold promise for letting us understand how different regions of the brain interact -- as in language behavior (Arbib and Caplan, 1979) -- when we seek to make contact with the level of analysis of the neurologist concerned with brain lesions, rather than that of the neurophysiologist and neuroanatomist who can trace a few cells at a time.

When we turn to the analysis of neural networks *per se*, we find that the 'continuous' style of analysis of cooperative computation is playing an increasing role. We have already mentioned that the formation of the underlying tissues has been studied in terms of reaction-diffusion equations. In the rest of this section, we briefly discuss schemes that explain how one part of the brain can be 'wired up' to another, analysis of the effects of early environment on the modification of feature detectors via synaptic plasticity, as well as models of cooperativity in the mature function of the nervous system. Certain of these studies will be taken up in more detail in subsequent sections.

By way of background from neurophysiology, we first recall the notion of retinotopy.

2.1 *Retinotopy and Population Encoding*: When light is focused by the lens upon the retina, a very small solid angle in the external world can affect each individual receptor. But as we move back through the layers of the retina and along the optic tract into different regions of the brain, neighbor neuron talks to neighbor neuron in affecting neurons further along the pathway. Thus, the activity of neurons away from the periphery may be influenced by 20 degrees or more of the visual field. Yet, as we move across layers of neurons in visual cortex, tectum (midbrain visual region) or lateral geniculate (thalamic visual region), we find that there is a lawful direction across the surface corresponding to up and down in the visual field and another corresponding to left and right. This property of preserving the spatial structure from the retina as we back up into the brain is called *retinotopy*.

Just as we may speak of retinotopy when we see a variety of 'maps' of the retina within the brain (Allman, 1977), so may we talk of somatotopy when we look at 'maps' of the body surface -- whether they be sensory maps of the tactile and other stimuli to the skin, or motor maps of the distribution of contraction of the musculature around the body (Brodal, 1969).

In such layered structures, it may be inappropriate to regard a single neuron as by itself conveying a vital message for the brain. Rather, it is the pattern of firing as distributed across a whole array of neurons that robustly encodes vital information about the world (Erickson, 1974). If this is so, the occasional misfiring of an individual neuron poses little problem because the receptive field that it samples overlaps the receptive field of hundreds or even thousands of other neurons in its vicinity in its own layer of the brain, yielding a natural redundancy and stability. This natural redundancy and stability, induced by the way in which layers of neurons represent spatial properties of the world, seems to solve the problem of "reliable computation in the presence of noise".

The 'maps' within a neural layer are not simple point-by-point transmissions of arrays of stimulation from the periphery. Rather, they involve sophisticated transformations. For example, Lettvin et al. (1959) identified several classes of ganglion cells in the retina of the frog, including cells which seemed most responsive to the presence of small wiggling objects in their receptive field and those that responded best to the passage of a large dark object across their larger receptive field. Moreover, they found that these different types of cells distributed their messages to the tectum in such a way that each cell-type projected to a different depth, with each projection retinotopic and corresponding points in the arrays one atop the other.

In the mammalian retina, Kuffler (1953) had found that the ganglion cells did not respond to these 'frog relevant' stimuli, but rather served to enhance contrast, while Hubel and Wiesel (1962) found 'edge detectors' in the visual cortex of cats and monkeys, which seemed to respond best to the movement of edges with a specific orientation in their receptive field. In this way, we find that the input to the brain is arrayed in maps of distinctive features. The suggestion is that these maps serve to provide the input variables for controlling the animal's behavior. In Section 7, "Modelling Frog Visuomotor Coordination", we shall develop this idea further in a specific situation. We here turn to a brief survey of cooperative computation in retinotopic arrays.

2.2 <u>Retino-tectal</u> <u>Connections</u>: We have seen that fibers from the retina reach the visual mid-brain, the tectum, and there form an orderly map. Sperry (1944) showed that this retinotopy of the tectum would, at least in the frog, survive rotation of the eyeball after section of the optic tract. After such an operation, nerve fibers growing out from the retina would still find the 'original' tectal locus. This might suggest that each fiber bears with it a unique 'address' and goes directly to the target point on the tectum. However, experiments in the 60's (reviewed in (Gaze, 1970)) showed, e.g., that if half a retina were allowed to innervate a tectum, then the map would expand to cover the whole tectum; while if a whole retina innervated half a tectum, then the map would be compressed. In other words, the fibers had in some sense to 'sort out' their relative position in using available space, rather than simply going for a prespecified target.

There are now a number of models which explain this phenomenon not in terms of an overall global organization principle, but rather by local interactions of a few fibers and the portion of tectum upon which they find themselves. These models include the arrow model of Hope, Hammond and Gaze (1976), the 'market' model of von der Malsburg and Willshaw (1977), and the branch-arrow marker model of Overton and Arbib (to appear). (See also Malsburg's chapter in this volume, which -- in common with the present section -- argues for the broader implications of such models for brain theory.)

2.3 Cortical Feature Detectors: Many cells in visual cortex are tuned as edge detectors, and it is known that the binocularity of this tuning can be modified by decoupling the input from the two eyes (Hubel and Wiesel, 1965). Moreover, early visual experience can drastically change the population of feature detectors (Hirsch and Spinelli, 1971; Blakemore and Cooper, 1970). This suggests that the growth of the visual system in the absence of patterned stimulation can at most sketch the feature detectors, and that interaction with the world is required whether to fully express the 'normal' situation, or to adapt an array of detectors to the peculiarities of a given environment.

These considerations have challenged a number of workers to come up with models in which there are mechanisms of local synaptic change within a neuron based on correlation between presynaptic and postsynaptic activity, or on reinforcement signals, or on a combination of the two (Amari, 1977; Grossberg, 1970; Grossberg and Levine, 1975; Kohonen and Oja, 1976; von der Malsburg, 1973). In particular, when local synaptic change is coupled with inhibitory interaction between neurons, a group of randomly connected model neurons can eventually differentiate amongst themselves to give edge detectors for distinct orientations. However, it must be noted that while certain cells in areas 17 and 18 of visual cortex become well tuned as simple cells, other cells in area 18 and 19 become well tuned as hypercomplex cells -- and complex cells arise in all three areas. It thus requires a more subtle theory than any to date to understand what it is about the precursive cellular geometry that provides preconditions for the different patterns of learning in the two populations.

Recent experiments (Spinelli and Jensen, 1979) show that early visual experience may actually increase the area of cortex allocated to features innervated by a given area of the periphery, contrary to a view that feature detector changes may simply involve atrophy of those neurons that are seldom active. Amari (1980) has recently modelled topographic organization of two nerve fields connected by modifiable excitatory connections; and has indeed proved that a part of the postsynaptic field which is frequently stimulated comes to be mapped on a large area of the postsynaptic field.

2.4 Optic Flow and Stereopsis: Given two frames of visual information, we may ask how local features in one frame are matched with the correct features in the other. If the two frames are taken in temporal succession, then the stimulus matching problem is that of computing optic flow; if the two frames come from two simultaneous vantage points (the left and right eyes, for example), then the stimulus matching problem is that of stereopsis.

There are now a wide variety of algorithms which show how one can start with an initial set of hypotheses about matching and then use local interactions to slowly change these hypotheses until obtaining a coherent segmentation of the image on the basis of common motion or common depth. While a number of models of stereopsis were based purely on cooperative computation (Dev, 1975; Julesz, 1971; Marr and Poggio, 1977b; Sperling, 1970), Marr and Poggio (1979) have offered an alternative model which uses channels of different spatial frequency to obtain a coarse match which can then be refined without cooperative computation. However, Frisby and Mayhew (1977) have offered psychophysical evidence for a modification of the latter model which does require cooperative computation between the channels of different frequencies. We shall provide a specific cooperative algorithm for the computation of optic flow in Section 5.

2.5 The Continuity of Development and Function: We have seen a mathematical commonality which linked reaction-diffusion models of basic pattern generation in biological tissues; the formation of projections from one region of brain to another as in the connection of retina to tectum; the tuning of connections to a cell within a tissue as in the formation of feature detectors in visual cortex; and the actual function of a brain region as in the computation of optic flow or stereopsis. This all suggests that many developments in brain theory will result from looking for a commonality of underlying mechanisms between neuroembryology, formation of connections, and adult function. We may summarize this with the slogan of (Arbib, 1972): "The brain is a somatotopically organized distributed layered computer." Nonetheless, it must not be thought that there is any single method to be 'plugged in' to solve all problems of brain theory. The remainder of this paper will illustrate the diversities as well as the unities.

3. INTERACTING SCHEMAS FOR MOTOR CONTROL

In this section we offer a top-down analysis of the interaction between perception, internal representations, and the control of movement. This will provide a setting for the bottom-up studies of Section 5, 6 and 7. We start by analyzing concepts from the theories of control and communication to be adapted to address the issues of perceptual-motor interaction.

Control theory has taught us how to break a system down into an array of continually active subsystems linked by message-bearing pathways through complicated patterns of feedback and feedforward. Computation theory has taught us how to break a complicated pattern down into a program describable by a flow diagram whose boxes correspond to the activation of various data transfers, tests and operations, and whose lines correspond to transfer of control from one box to another. In control, then, we have continually active systems in constant intercommunication; in serial computation we have activation of one subsystem after another, with the pattern of activation delicately determined by tests of current data. In this section, I want to suggest how 'programs in the brain' might be viewed as sharing properties of both control block diagrams and computer flow diagrams. The resultant approach involves concurrently active systems, with which systems are active at any time being determined on the basis of current interactions. I shall use the term 'schema' for the units of control from which these programs are built.

3.1 The Action-Perception Cycle: The notion of the action-perception cycle (Arbib, 1981a; Neisser, 1976) serves to emphasize that the current situation in the environment does not in general serve to determine the behavior of the animal. Feeling hungry, we go to the kitchen to get food from the refrigerator. Our brain models the world (Craik, 1943). Again, many movements are explorations of the world around us. In short, we perceive so that we may plan our actions appropriately; but in acting we provide ourselves with new opportunities to perceive. The cycle of action and perception continues.

3.2 The Schema-Assemblage: How is it that looking around us we come to recognize objects and their spatial relationships? We shall look at two theories of this process in Section 4. Here I want to stress that it seems reasonable to posit that we perceive the world on the basis of our own prior knowledge. We have 'schemas' which are programs or systems which let us recognize a phone or a person or a mountain from appropriate visual or other stimuli. I suggest that our knowledge of the world is divided into a short-term model that represents our appreciation of our current place in space and time, and a long-term model which represents all that we know both consciously and unconsciously. To a first approximation, and especially insofar as it refers to the sensible environment in which we currently find ourselves, the short-term model is then to be construed as an assemblage of activated schemas whose pattern of activations is related to the current state of the environment. By contrast, long-term memory is the distillation of experience (some of it genetic) represented by the repertoire of schemas available for activation. Of course, this begs many questions as to how schemas are coded in neural terms, and how they are related one to the other. (I stress that this is a top-down analysis, and that much remains to be done to determine to what extent 'schemas' and other processes described here are directly instantiated in neural processes.) Clearly, retinotopy provides a fine framework for understanding the disposition of those schemas which represent objects which are within our current central and peripheral visual field. But when we come to representing our current awareness of objects in other parts of the house or in the environment beyond, then we come to deep problems which I shall not enter into in the present paper.

3.3 Perceptual and Motor Schemas: The present theory posits that a schema is like a neural network, and may thus be in a state of greater or lesser activity. It is perceptual to the extent that its activation can vary on the basis of cues from both peripheral stimuli and internal context. It is motor to the extent that, being activated, it can determine an appropriate course of action. Just as a control system may have an identification procedure (cf. Section 6) to tune its parameters to provide appropriate control signals to that which it controls, so will in general a motor schema need to be linked to a perceptual schema so that the interaction with the environment is based on a proper appreciation of the nature of that environment. Thus, a schema for grasping an object can be guided by visual perception of not only the position of the object but also its size and orientation. A motor schema, in this sense, is akin to what the Russian school, founded by Bernstein (1968), has called a synergy (a different usage of the term from that of Sherrington).

3.4 Coordinated Control Programs: We now offer two examples of the notion of a coordinated control program which is to combine the features of both conventional control block diagrams and computer flow diagrams.

Figure 1 offers a hypothetical diagram of a possible program for the control of reaching towards a visually sensed object. At the top of the diagram we see the operation of the perceptual schemas which recognize the object and locate it on the basis of a description, and those perceptual schemas which in turn determine the size and orientation of the object. It has been observed by Jeannerod and Biguer (1980) that when a person reaches for an object he already begins to shape his hand to the size and orientation of the object during the reaching movement. Thus, in the motor schemas, we see the dashed line indicating simultaneous activation of both reaching and shaping of the hand. If we regard the shaping of the hand as part of the grasp schema, we have the interesting fact that the completion of this subtask does not initiate the next task of grasping. It must, in fact, wait for an activation -- possibly based upon feeling the object touch the palm of the hand -- consequent upon the successful completion of the reaching movement.

Figure 2 is based on Brinkman and Kuypers (1972) and Haaxma and Kuypers (1974). They were able to show that finely coordinated visually-guided behavior involved the cooperative computation of two different systems. A pathway involving the brainstem controls the undifferentiated hand movements akin to the simple grasping discussed in the previous example. A pathway from visual cortex to precentral gyrus and thence directly via the pyramidal tract to motor neurons controls the distal musculature, and is responsible for the control of relatively independent finger movements. With interruption of either the corticocortical connections or of the pyramidal tract, the animal was unable to shape its hand in such a way as to dislodge a pellet from a groove whose orientation could be visually determined. Instead, the animal could reach for the pellet but without preshaping the hand, and would then move its hand back and forth under tactile control until by chance the pellet was dislodged. At that time, the tactile feedback sufficed to allow the animal to grasp the pellet efficiently and bring it to his mouth.

We shall place these examples in perspective when we relate such top-down models to bottom-up models in Section 8.

4. TWO THEORIES OF VISION

Section 3 suggested that a task of perceptual systems was to update an assemblage of representations of particular objects in spatial relationship. This problem is of intense interest to many workers in the branch of artificial intelligence known as 'machine vision' or 'computer vision', irrespective of any question of the use to be made of the representation of visual input. This section examines two approaches to machine vision, those of Hanson and Riseman (1978a) and of Marr (1978), and considers to what extent they can be viewed as cognitive models, rather than simply as programs to get machines to emulate certain aspects of vision. In fact, Hanson and Riseman view themselves as working primarily at designing a computer system <u>per se</u> to do something useful, while Marr does see his approach as providing a top-down specification of the visual system whether it be in brain or machine. However, I shall argue that the Hanson-Riseman approach provides valuable cues for brain theory which supplement, and occasionally challenge, those offered by Marr.

Before going on, I should note that in the rest of this section we shall be using the terms 'bottom-up' and 'top-down' in a somewhat different sense from our analysis of styles of brain theory. Within a visual system, bottom-up processes are those that proceed by more and more elaborate processing of the peripheral signal without reference to knowledge of objects in the world; while 'top-down' processes are those that proceed from knowledge of objects to an interpretation of 'lower-level' patterns. Both the models discussed here are in agreement that there must be two stages of bottom-up processing. The first stage applies a variety of local processors to come up with a map which highlights what are likely to be the most 'information bearing' places within the visual input. They also agree that a second stage must build upon this initial 'feature map' -- what Marr calls the 'primal sketch' -- to yield a representation which provides valuable information about the position and shape of objects without yet calling upon any interpretation based on knowledge of objects that might be in the world. Here, the two models depart quite drastically, with Hanson and Riseman advocating a pattern of segmentation, while Marr advocates the 2-1/2 D sketch. Having obtained these intermediate representations, both theories invoke high-level information to come up with hypotheses about what objects in the world could be responsible for the observed visual pattern.

The general scheme, then, is bottom-up processing through several levels of representation until 'world knowledge' can be invoked to generate hypotheses; these hypotheses then act top down to verify or disqualify themselves by determining whether or not other data from the visual image are compatible with those data which evoked the hypothesis in the first place. It should be noted that this same overall system organization can be seen in AI studies of speech understanding, as in the HEARSAY system (Erman and Lesser, 1980).

4.1 The Hanson-Riseman Approach: The Hanson-Riseman approach uses a pattern of segmentation as its intermediate representation. An attempt is made to segment the image into different regions. There are two methods of segmentation. One is 'edge-finding' based on discontinuities in color or texture or depth which could signal a break between two surfaces. The other is 'region-growing', aggregating areas of similar visual stimulus by finding clusterings in the feature space and then mapping representative symbolic labels back upon the image to determine a partition of the visual field. The variety of shapes and illuminations in the world are such that it proves virtually impossible, at least with current techniques, to come up with regions that are in 1:1 correspondence with surfaces of distinct objects. The pattern of light and shade in a tree can break it into a number of chromatically distinct regions. A highlight may make it impossible to see an edge separating one region from another. Shadows and highlights may themselves be treated as distinct regions rather than features lying upon a given surface. In fact, there is a hierarchical problem of grouping texture elements -- consider leaves, clumps of leaves and branches, trees on a hillside, etc. While it is true that more sophisticated bottom-up processing can be designed which will take into account various processes of color change under highlighting and shadowing to allow merging of regions which would be separated on a crude analysis, it nonetheless seems fair to posit that total segmentation cannot be done without invoking real-world knowledge. Note, however, that segmentation can be improved by a process of competitive cooperation between different segmentation processes, such as those based on edge-finding and those based on region-growing. Once this approximate segmentation is completed, the Hanson-Riseman approach calls various processes into play to make hypotheses about the objects of which those regions are parts of surfaces. The 'schemas' which represent objects or other visual regions must thus contain the necessary routines to determine whether that which they represent is present within the scene. For example, if a region is blue and is near the top of a picture taken outdoors with a level camera, then a reasonable hypothesis is that the region is sky. If it is green and near the bottom of a picture in an outdoor scene, then it is a plausible hypothesis that it is grass. Other cues may suggest the presence of bushes, trees, houses, windows, cars, etc. Such hypotheses can then be checked by seeing whether the region can be merged with other regions which satisfy bounds on their spectral attributes and other features, and if the resulting posited surfaces are appropriate in terms of shape, with extra processing required to take occlusion effects into account, and size. A camera model can be invoked to infer the size of an object from an estimate of its distance along the ground plane in the case where other depth information from stereo disparity or a range finder, say, is unavailable.

A three-dimensional object is represented in the Hanson-Riseman system by the numerical representation of surface patches which approximate the surface surrounding the volume of an object. The patch boundaries are coded by the use of cubic splines to yield a representation akin to that used in computer aided design of cars (York, 1980)

In the top-down matching of a scene against a posited schema, a 2-D schema is generated corresponding to the view of the 3-D schema from a particular perspective. To increase flexibility, the 2-D representation is not a picture with fixed edges, but is rather a graphical structure which posits a range of likely positions for objects with corresponding flexibility in the specification of edges. Once an object is posited in a particular region, the system can check the hypothesis by seeing if in a particular area of the image, top-down processes can merge regions from the segmentation to meet the spectral shape and size characteristics that can confirm the original hypothesis. However, while the current system contains routines to confirm certain hypotheses (Parma et al., 1980), the design of programs for inferring hypotheses from the patch representations is still at an early stage of development.

4.2 The Marr Theory: In Marr's theory, the intermediate representation after the 'primal sketch' is not based on segmentation, but is based upon a depth analysis which seeks to assign to each point of the scene an estimate of its orientation in space (recall the review of stereopsis in Section 2; other cues come from surface highlights (Horn, 1974) and motion (Ullman, 1979)). The resultant 2-1/2 D sketch is like a bas-relief, in which the shape is determined at each point, but there is no symbolic representation of separation into distinct objects. Further computations are then designed to find axes of symmetry for different bulges in the 2-1/2 D sketch, and the resultant collection of axes are to be used as stick figures which provide access to a data base of different objects known to the system. In other words, the three-dimensional representation here is not based on surface patches but on 'body-centered coordinates' of the represented object, providing, as I have said, a sort of stick figure skeleton, with a specification of cross sections to be swept up and down these axes to 'flesh out' the three-dimensional object. Again, as in the Hanson-Riseman system, programs for such data-base access remain a topic for future research, and I believe that it is open to question whether in fact the 2-1/2 D sketch can be constructed with sufficient accuracy to drive the process of axis inference well enough to allow reliable retrieval of hypotheses.

The two approaches as currently constituted provide a useful base for understanding the visual system, but certainly do not stand alone. Data are often noisy, so that inference of the 2-1/2 D sketch, or the segmentation sketch, is unreliable, and can at best suggest hypotheses rather than lead to the selection of a unique hypothesis. Again, a strategy of cooperative computation between depth estimation and segmentation might well work better than either process alone, and computer experiments along this line would seem well worth while. Marr's approach posits a uniquely three-dimensional representation of objects, rather than accepting the perhaps more plausible view (cf. Minsky's theory of frames (1975)) that our knowledge of an object is often a synthesis built upon our viewing of it from a number of different perspectives. Finally, while Marr has downplayed cooperative computation in his latest stereopsis algorithm, I believe, as already suggested twice above, that the proper development of a theory of vision systems, synthesizing and building upon features of many different approaches (sampled in (Hanson and Riseman, 1978b)) will involve 'cooperative computation' between a multitude of processes: In a feature-rich environment, there are always more features available than can be taken into account in a reasonable time. It is thus necessary for processes to be initiated which extract certain salient features; but the system must be so designed that the use of these features does not preclude taking into account other features. We saw that in the Hanson-Riseman approach, a process initiated on the basis of feature measure cues could then be rigorously checked by invoking other processes which could take size or shape into account. This type of interaction of multiple knowledge sources is, it seems to me, the style of the brain with its incessant interaction of hundreds of continuously active brain regions.

5. COMPUTING THE OPTIC FLOW

We have seen that machine vision research postulates the high-level systems to build upon the representations initially determined at the low level (as in the 'primal sketch') to utilize 'perceptual schemas' to recognize objects within the environment. J. J. Gibson (1955; 1966; 1977) was one of the people who most forcefully made clear to psychologists that there was a great deal of information that could be 'picked up' by 'low-level systems' and that, moreover, this information could be of great use to an animal or to an organism even without invocation of 'high-level processes' of object recognition. For example, if, as we walk forward, we recognize that a tree appears to be getting bigger, we can infer that the tree is getting closer. What Gibson emphasized, and others such as Lee (1974; Lee and Lishman, 1977) have since developed, is that it does not need object recognition to make such inferences. In particular, the 'optic flow' -- the vector field representing the velocity on the retina of points corresponding to particular points out there in the environment -- is rich enough to support the inference of where collisions may occur within the environment and, moreover, the time until contact. We shall detail elsewhere (Arbib et al., to appear) our current studies of how inference from the optic flow might be used in directing the locomotion of a robot around an obstacle-cluttered world, and the subsequent analysis of how relevant such study may be to the cognitive psychologist. Here, I want to emphasize a problem often glossed over in Gibson's writings, namely that of the actual computation of the optic flow from the changing retinal input. Our studies to date have been 'in the style of the brain', but have not been related to actual neural circuitry. In what follows, then, rather than asking how neurons might pick up the optic flow on the basis of continuously changing retinal input, we shall simply offer an algorithm (Prager (1979), Prager and Arbib (to appear); for other algorithms, see, e.g., Ullman (1979)), played out over a number of interacting layers each of which involves parallel interaction of local processes, where the retinal input is in the form of two successive 'snapshots', and the problem is to match up corresponding features in these two frames. (Mathematically, then, the problem is the same as that of stereopsis, as discussed in Section 2. However, whereas there are only two eyes, there may be many successive moments in time, so that the initial algorithm for matching a successive pair of frames can be vastly improved when the cumulative effect of a whole sequence can be exploited.)

The problem is posed in Figure 3, where we see four features extracted from Frame 1, shown as circles, and four features from Frame 2, represented as crosses. The stimulus-matching problem is to try to match up the features in the two frames that correspond to a single feature in the external world. Figure 3(a) shows an assignment that seems far less likely to be correct than that shown in Figure 3(b). The reason that we would, lacking other information, prefer the latter stimulus-matching is that the world tends to be made up of surfaces, with nearby points on the same surface being displaced similar amounts. (This use of the plausible hypothesis that our visual world is made up of relatively few connected regions to drive a stimulus-matching process was enunciated, for stereopsis, in (Arbib et al., 1974).) Our algorithm, then, will make use of two consistency conditions:

FEATURE MATCHING: Where possible, the optic flow vector attached to a feature in Frame 1 will come close to bringing it in correspondence with a similar feature in Frame 2.

LOCAL SMOOTHNESS: Since nearby features will tend to be projections of points on the same surface, their optic flow vectors should be similar.

In developing an algorithm which is 'in the style of the brain', we shall assume that there is a retinotopic array of local processors, which can make initial estimates of the local optic flow, but will then pass messages back and forth to their neighbors in an iterative process to converge eventually upon a global estimate of the flow. The need for interactions if a correct global estimate is to be obtained is shown in Figure 4, where we see a local receptive field for which the most plausible estimate of the optic flow is greatly at variance with the correct global pattern. Our algorithm is then as shown in Figure 5. We fix two frames, and seek to solve the matching problem for them. An intial assignment of optic flow vectors might be made simply on the basis of nearest match. The algorithm then proceeds through successive iterations, with the local estimate for the optic flow vector assigned to each feature of Frame 1 being updated at each iteration. (Computer simulations suggest that at most 20 iterations of the algorithm are required to yield convergence to a reasonable global estimate of the flow.) Consider, for example, the Frame 1 feature A of Figure 5, and the position B which is the current hypothesis as to the location of the matching stimulus in Frame 2. We see that, were feature matching to be the sole criterion, the new optic flow would be given by the wavy arrow which matches A to the feature in Frame 2 closest to the prior estimate, namely B. On the other hand, if only local smoothness were taken into account, the new optic flow vector assigned to A would be the average of the optic flow vectors of features within a certain neighborhood. Our algorithm updates the estimate at each iteration by making the new optic flow estimate a linear combination of the feature matching update and the local smoothness update, as indicated by the dashed arrow emanating from A in Figure 5. As stated, the algorithm works quite well in giving a reliable estimate of optic flow within 20 iterations.

If we take advantage of the availability of a whole sequence of frames, rather than just two, then we can obtain both an increasingly accurate estimate of the optic flow, and with less iterations to handle each new frame as it is introduced. For example, if, having matched Frame n to Frame n+1 we try to match Frame n+1 to n+2, it is reasonable to assume that -- to a first approximation -- the optic flow advances a feature by roughly the same amount in the two frames. If we thus use the repetition of the previous displacement, rather than a nearest neighbor match, to initialize the optic flow computation of the two new frames, we find from simulations that only 4 or 5 iterations, rather than the original 20, are required, and that the quality of the match on real images is definitely improved.

The algorithm just described is based on two consistency conditions, feature matching and local smoothness. It is instructive to note where these constraints break down. If one object is moving in front of another object then points on the rear surface will either be occluded or disoccluded during this movement, depending on whether the front object is tending to cover or uncover the object behind it. Thus, if we look at the current estimate of the optic flow and find places where the flow vector does not terminate near a similar feature to that from which it starts, then we have a good indication of an occluding edge. On the other hand, the local smoothness will also break down at an edge, for the two objects on either side of the edge will in general be moving differentially with respect to the organism. Thus, we can design edge-finding algorithms which can actually use the breakdown of our consistency conditions to find edges in two different ways, on the basis of occlusion/disocclusion, and on the basis of optic flow discontinuity. To the extent that the estimate of edges by these two processes is consistent, we have the cooperative determination of surfaces within the image. What is interesting is that, to the extent that good edge estimates become available, the original basic algorithm can be refined, as shown in Figure 6. (This extension of the algorithm has not yet been implemented.) For now, instead of having 'bleeding' across edges, we can dynamically change the neighborhood of a point, so that the matching of features or the conformity with neighboring flow can be based almost entirely upon features on the same side of the currently hypothesized boundary. (But not completely, for at any time the edges will themselves be confirmed with limited confidence, and so may be subject to later change.)

We thus see in Figure 6 an 'evolutionary design process'. The basic algorithm (1) provides new information which can then be exploited in the design of the cooperative segmentation algorithms (2), but once the segmentation information is available, the original algorithm can be refined by the introduction of segmentation dependent neighborhoods (3). I suggest that this is not simply an interesting engineering observation, but gives us some very real insight into the evolution of the brain: basic systems provide the substrate upon which 'higher level' systems may evolve; but these higher level systems then enrich the environment of the lower systems, and these lower level systems may then evolve to exploit the new sources of information. While it is still useful, to a first approximation, to talk of low-level and high-level systems, we see that there is no longer any univocal flow of information. We are very close to the Jacksonian notion of levels (Jackson, 1874; 1878).

6. A MODEL OF THE CEREBELLUM

To see how the top-down analysis of Section 3 can make contact with details of neural circuitry, we now examine a model of the cerebellum (Arbib et al., 1974; Boylls, 1975; 1976). The model brings together the notion of a motor schema with the notion of maps as control surfaces, and is important in that it exhibits neural layers acting as control surfaces representing levels of activation for the coordination of muscles, complementing our study of retinotopic representations of visual input.

We suggested in Section 3 that the problem of motor control is one of sequencing and coordinating motor schemas, rather than directly controlling the vast number of degrees of freedom offered by the independent activity of all the motor units. We must not only activate the appropriate schemas but must 'tune' them. To understand this notion of tuning we need an important concept from modern control theory, that of the identificaiton algorithm. In the familiar realm of feedback control theory, a controller compares feedback signals from the controlled system with a statement of the desired performance of the system to determine control signals which will move the controlled system into ever greater conformity with the given plan. The appropriate choice of control signal must depend upon having a reasonably accurate model of the controlled system -- for example, the appropriate thrust to apply must depend upon an estimate of the weight of the object that is to be moved. However, there are many cases in which the controlled system will change over time in such a way that no a priori estimate of the system's parameters can be reliably made. To that end, it is a useful practice to interpose an 'identification algorithm' which can update a parametric description of the controlled system in such a way that the observed response of the system to its control signals comes into greater and greater conformity with that projected on the basis of the parametric description. When a controller is equipped with an identification algorithm, and when the controlled system is of the class whose parameters the algorithm is designed to identify, and when, finally, the changes in parameters of the controlled system are not too rapid, then in fact the combination of controller and identification algorithm provides an adaptive control system, which is able to function effectively despite continual changes in the environment.

So far, our analysis has been top-down. We now turn to the neurophysiological data. As a 'working example' we shall model cerebellar function in locomotion of the high decerebrate cat (Shik et al., 1966). Where Sherrington had noticed that stimulation of Deiters nucleus in the standing animal would lead to extension of all the limbs, Orlovskii found that in the high decerebrate cat, stimulation of Deiters nucleus during locomotion would not affect extension during the swing phase, but would yield increased extension during the support phase. Since the locomotory 'motor schema' has been shown to be available even in the spinal cat in (both classical work by Sherrington (1910) and modern studies (Herman et al., 1976)), it seems reasonable to view the system in which the cerebellum and Deiters nucleus are involved as providing an identification algorithm for the parametric adjustment of the spinal schema. We now turn to Boylls' model which shows how the adjustment of those parameters might be computed within the cerebellar environs.

As is well known (Eccles et al., 1967), the only output of the cerebellar cortex is provided by the Purkinje cells, which provide inhibitory input to the cerebellar nuclei. Each Purkinje cell has two input systems. One input is via a single climbing fiber which ramifies and synapses all over the Purkinje cell's dendritic tree. The other input system is via the mossy fibers, which activate granule cells whose axons rise up into the layer of Purkinje cell dendrites (which are flat, with the planes of all their dendritic trees parallel to one another) to form T's, whose crossbars run parallel to one another at right angles to the planes of the Purkinje dendritic trees. (There are a number of interneurons in the cerebellar cortex, but we shall not model these here, but shall instead concentrate on the basic cerebellar circuit of mossy and climbing fibers, and of granule and Purkinje cells.)

The climbing fiber input to a Purkinje cell is so strong that, when its climbing fiber is fired, a Purkinje cell responds with a sharp burst of four or five spikes, known as the climbing fiber response (CFR). Many authors have thought that the 'secret' of the climbing fiber is this sharp series of bursts, but we shall suggest below that the true role of the climbing fiber input is to provide the suppression of Purkinje cell activity for as much as 100 milliseconds which has been found to follow the CFR (Murphy and Sabah, 1970).

The overall architecture of Boylls' model as played over an array of interacting control surfaces is shown in Figure 7, which is an anatomical template of circuitry ubiquitous in cerebellar transactions. That is, specific labels could be given to, say, the 'brainstem output nucleus' as red or Deiters nucleus, the 'reticular nucleus' could be reticularis tegmenti pontis or paramedian, etc. From this architecture we gather that the output from the cerebellar nuclei via the brainstem 'output' nucleus results from the interaction between cerebellar cortical inhibition as supplied by the Purkinje cells and between drives from the reticular nucleus. Tsukahara (1972) has demonstrated the possibility of intense reverberation between the reticular and cerebellar nuclei following removal of Purkinje inhibition, and Brodal and Szikla (1972) and others have demonstrated the anatomical substrate for such loops, with a somatotopy as indicated in Figure 8. We thus postulate that there will be explosively excitatory driving of the cerebellar nucleus by reticulo-cerebellar reverberation unless blocked by Purkinje inhibition.

The output of cerebellar tuning is expressed as a spatio-temporal neuronal activity pattern in a cerebellar nucleus, which can then be played out via the brainstem nuclei to spinal levels. A careful analysis of the anatomy enabled Boylls to predict that the agonists of a motor schema would be 'represented' along a saggital strip of the cerebellar cortex, while its antagonists will lie orthogonal to that strip (in the medio-lateral plane). Applications of this formula to cortical topography of the anterior lobe, as developed by Voogdt (1969) and Oscarsson (1973), allowed Boylls to identify particular cortical regions as associated with equally particular types of hindlimb-forelimb, flexor-extensor synergic groupings. This led to conclusions which are experimentally testable.

The Boylls model suggests that activity within the cerebellar nucleus is initiated through topically precise climbing fiber activity; the mechanism involves their direct cerebellar nuclear activation coupled with the suppression of the target Purkinje cell activity in the cortex via the above-mentioned 'inactivation response'. Once activity is installed in cortico-nuclear interactions via climbing fiber intervention, the underlying reverberatory excitation helps to retain or 'store' it. At the same time, this activity is transmitted to the cerebellar cortex on mossy fibers, eventually altering the inhibitory pattern in the nuclear region surrounding the active locus. The relevant pathways involving mossy fibers and the cortico-nuclear projection are schematized in Figure 9. The spread of parallel fibers indicated in Figure 9 (for simplicity of the diagram, the granule cells are omitted) yields a form of lateral inhibition which provides spatial 'sculpting' in a way depending on the elaborate geometry of cerebellar cortex and cortico-nuclear projections. Mossy inputs of various types tune the resultant patterns to the demand of the periphery; and the program is spinally 'read out' as appropriate.

Testing of the various hypotheses has required computer simulation of this neuronal apparatus. Simulation results corroborated the conjecture that cerebellar related circuitry could support the short-term storage of motor schema parameters initiated (and periodically refreshed) by climbing fiber activity. Figure 10 suggests a typical nuclear activation pattern so introduced.

7. MODELLING FROG VISUOMOTOR COORDINATION

In my own group, we have chosen visuomotor coordination in frog and toad as a setting where the top-down and bottom-up approaches to brain theory may interact. Our initial model building was conducted at Stanford in 1970; current activity is in collaboration with a number of experimentalists including Ingle of Brandeis, Ewert of Kassel, and Collett of Sussex. After discussing the basic models of prey selection, and of prey-predator discrimination, we shall discuss how new behavioral experiments are being used to develop a top-down analysis which, we hope, will allow us to extend our current modelling of the animal's behavior to models which take context more fully into account.

Lettvin, Maturana, McCulloch and Pitts (1959) asked "What does the Frog's Eye Tell the Frog's Brain?" Didday (1970; 1976) and I asked "What does the Frog's Eye Tell the Frog?" It is one thing to say that the human monitoring a cell through a microelectrode can correlate the cell's activity with some feature of the external world; it is quite another thing to say that the neural circuitry within the brain of the animal can actually make use of that information in determining behavior (Perkel and Bullock, 1968). We sought to ask, then, how the 'bug detector' information from the retina might be used to guide the animal's activity. Our basic perspective was formed by the behavioral studies of David Ingle (1976), who noted that the frog would snap -- orient and zap with the tongue -- at a suitably small wiggling stimulus within a certain range about the animal. Moreover, if two stimuli were presented within the snapping range, while the animal would often snap at just one of them, there were cases in which it would snap at neither, or would snap in between at the 'average fly'.

'Pure' top-down analysis of prey selection would specify the task as, e.g., "Develop a procedure for finding the greatest element in an array of elements." Unless constrained by the requirement "use local, parallel computations," this might be realized on a computer by simply scanning a list of values to find the maximum; and the 'no fly' and 'average fly' effects could be handled by subroutines that would detect when the top two entries of the list bore some designated relationship. However, this serial process would not be interesting as a brain model whatever its utility as a computational summary of the behavior. We thus asked how this process of selection could be played out through the interaction of neurons rather than through the supervention of some executive program. The model that we finally developed involved an array of neurons being modulated by another array of inhibitory neurons in such a way that peaks of activity in the first layer would compete through the second layer. In general, the highest peak on the first layer would finally suppress all other peaks and emerge from the system to control motor activity. However, in some cases two peaks would be sufficiently well matched to hold a balance in which each held the other below the threshold for action.

Recent advances in neurophysiology have led to better identification of different cell types, and the original model which was thought to correspond to cells just in the tectum is now conceived of as resulting from interaction of cells in tectum and pretectum (Figure 11). In particular, we are now looking at the detailed anatomy of the tectum to model it as an array of 'tectal columns' (Lara, 1980; Lara et al., to appear). Ingle (1975) observed that presentation of a fly-like stimulus to a frog for 0.3 seconds will rarely elicit a response, presentation for 0.6 of a second usually will. If a stimulus is presented for 0.3 of a second, taken away for several seconds and then presented again for 0.3 of a second, then this second short presentation will be enough to elicit a response. We have modelled this short-term memory. We have now shown that a plausible geometry of the tectal column will yield this facilitation effect.

In the model of prey selection, we have asked what processes can modulate an array of activity on the tectal map which serves as control surface for the snapping response. We have thus moved one step beyond a simple stimulus-response model, in that we are now asking how given a structured stimulus one part of it can be selected for response. The study of facilitation also shows that the response is 'history dependent'. Ewert (1976) has addressed a different question. If an object is small and moves in a certain way it is a prey to be snapped at; if it is large and moves in a certain way it is a predator to be jumped away from. How does the animal make this discrimination? Ewert has shown that if the pretectum-thalamus region of the toad (which is similar enough to frog for the present study) is ablated, the animal will snap at any moving object, no matter how large it is. This has suggested that the tectum is not, as we at first thought, a device to guide snapping at flies, but is rather a device to guide snapping at moving objects. It is then the task of other 'higher' brain regions to modulate this tectal activity, 'differentiating' the process of recognition to make finer discriminations. Working with von Seelen, Ewert developed a model (1974) of filters in both tectum and pretectum of toad with inhibitory modulation of tectal activity when pretectal cells signalled the presence of 'anti-worms'. One of our current efforts is to synthesize -- in a way consistent with current behavioral, anatomical, and physiological data -- our original model of prey selection with Ewert's work on prey-predator discrimination.

The above models of frog and toad treat the environment as made up of a number of objects. The animal in essence has to choose which of these objects to respond to and whether to respond to that object as prey or predator. We now briefly discuss a number of experiments which bring 'context' into the world of the animal. Our group has outlined models for them of an artificial intelligence kind -- they are programs, rather than neural nets which will yield the behavior. Nonetheless, unlike the serial list approach to prey selection mentioned briefly above, they are structured in terms of interacting processes, so that they can be used as plausible hypotheses about the interaction between brain regions. This is in the spirit of many workers in AI who use psychological observations to design their models of how to achieve some intelligent behavior, and structure the models in terms of concurrently active, interacting processes. However, lesion studies further constrain the model neurally. We thus hope to see in the near future not only the fuller articulation of the top-down model on the basis of recent behavioral experiments, but also physiological modelling and experimentation.

Ingle (1976) observed that when confronted with a 'predator stimulus', the direction of escape of a frog would be a compromise between the forward direction and the direction directly away from the stimulus. When he interposed a barrier in the preferred escape path of the frog, he found that the preferred path was no longer taken, but the frog would tend to jump to one side or other of the barrier. A plausible interpretation, then, is that the animal would come as close to the preferred route as it could while avoiding an obstacle in the way. Here, then, we see context at work: the animal is no longer simply responding to the aversive stimulus, but it is rather integrating the spatial structure of the world around it in choosing its response.

In more recent work, Collett (1977; Lock and Collett, 1980) has observed not only the animal's behavior with respect to barriers, but has also looked at toads faced with a chasm between them and a worm. He finds that if the chasm is narrow enough the animal will jump across, if it is wide but shallow the animal will step down and walk across, but if the chasm is both deep and wide the toad will simply turn away.

In summary, then, we see that in the study of visuomotor coordination in frog and toad we have a preparation in which we can both carry out subtle behavioral studies to bring 'manageable' cognitive aspects into the analysis; while at the same time we find interesting subproblems where detailed neural circuit analysis is possible.

8. A STYLE FOR THE BRAIN

The very richness of current research on the brain guarantees that a view such as this must be incomplete. Most neuroscience is experimental, and emphasizes the cellular and subcellular levels, with especial emphasis on chemical mechanisms. By contrast, while stressing the need for interaction with experiments, the present paper focuses on theory, with especial emphasis on models which relate neural models ('bottom-up') to the analysis of cognitive processes in terms of interacting function ('top-down'). Even here, our examples of cognitive processes stress vision and the control of movement (see (Arbib, 1981b; 1981c; Arbib and Caplan, 1979; Lavorel and Arbib, to appear) for a parallel attempt to relate cooperative computation to clinical studies of brain and language). Within this chosen area, we have sampled a number of fruitful models, with varying degrees of neural veridicality -- optic flow, cerebellar tuning of motor schemas, visuomotor coordination in frog and toad Below, I shall try to extract a general lesson from such models. But first, I want to return to the top-down analysis of Section 3 to put these models into perspective. In the language of that section, we may say that we begin to have neural models of individual 'schemas', but very little understanding of the way in which 'coordinated control programs' are constructed, embodied, or executed in neural tissue. Studies of neural circuitry in breathing and rhythmic locomotion (cf. the companion article by Stent) provide a useful first sequence of steps, but do not address the problems of coordinated phasing in and out of diverse activities in a goal-dependent way. The awareness of such shortcomings is not intended as a criticism of the current state of our science, but rather as a spur to the further articulation of top-down analyses to better 'stretch' the scope of neural modelling.

Given this perspective, the rest of the paper is devoted to a summing-up of a view of cooperative computation as "a style for the brain". Here, I do not wish to reiterate the case for a cooperative analysis of neural interactions in layers of neurons, whether in development, learning, or perceptual-motor function -- a case made in Section 2, and argued for persuasively by Christoph von der Malsburg in his discussion of the ontogenesis of retinotopy as a paradigm of organization in the brain. Nor do I want to say more of the 'discrete' form of cooperative computation as embodied in such AI systems as VISIONS and HEARSAY. Rather, I want to recall an intermediate level of cooperative computation which seems to give us general insight into layer-by-layer interactions within the brain.

The optic flow algorithm of Section 5 first developed a basic system of cooperative computation of optic flow within a retinotopic array, and then developed high-level systems to conduct segmentation on the basis of the output of this array; the higher-level systems then enrich the environment of the basic systems, which can then be adjusted to exploit the new sources of information. The distinction between low and high levels becomes blurred -- the 'low-level' systems provide the necessary data to initiate 'high-level' hypotheses, but these hypotheses are needed to constrain the lower-level processes. This is very much the approach to visual systems espoused in Section 4. There is no longer a simple one-way flow of information, but rather a 'coming to equilibrium' of multiple systems -- an equilibrium which is, of course, dynamic as the 'retinal' input to the system changes with time. The system must possess something of the 'adiabatic approximation' of Haken (1978, Chapter 7) -- it must be able to adjust to significant changes in the world at a faster rate than that at which those changes occur. In the same fashion, the interacting layers of the cerebellar model of Section 6 must adjust the tuning parameters of the motor schemas rapidly enough to keep pace with changes in the environment -- this is what makes the system adaptive, with perceptual processes (identification) intimately intertwined with control.

Again, we saw that the tectum of frog and toad could be regarded as a basic array for the control of snapping movements, while the pretectum apparently serves to differentiate the recognition of moving objects so that snapping would not be directed to large objects. In the same way, the work of Kuypers et al. (recall Figure 2 of Section 3) stresses that we should not view the pyramidal and extra-pyramidal pathways as alternative paths for motor control. Rather, the pyramidal pathway serves to differentiate and refine distal movements which ride atop movements of proximal musculature under extrapyramidal control.

Interestingly, making contact between frog and monkey studies, it appears that the monkey's basic 'undifferentiated' reaching movements can work quite well with visual input from the retina directed through the superior colliculus (the mammalian analog of the tectum) even in the absence of visual cortex. In fact, primates can exercise quite elaborate control of their movement in the absence of visual cortex. In his study of "What the Frog's Eye tells the Monkey's Brain," Humphrey (1970) built on the argument that a monkey without visual cortex should have at least as much visual ability as a frog. But monkeys had hitherto appeared to be blind when they had lost visual cortex. Humphrey argued that the monkeys had not been taught to pay attention to the visual cues they have, and after two years he was able to get a monkey without visual cortex to be able to grab at moving objects, and use changes in luminance -- such as an open door -- for navigation, even though delicate processes of pattern recognition were never regained.

In this way, we see the development of a paradigm which can guide us in the mathematical analysis and computer simulation of interacting brain regions as we come to address an ever richer array of cognitive processes.

REFERENCES

Allman, J., 1977, Evolution of the visual system in the early primates, in "Progress in Psychobiology and Physiological Psychology", pp. 1-53, Academic Press, New York.

Amari, S., 1977, Neural theory of association and concept-formation, Biol. Cybernet. 26: 77-87.

Amari, S., 1980, Topographic organization of nerve fields, Bull. Math. Biol. 42: 339-364.

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

Arbib, M. A., 1964, "Brains, Machines and Mathematics," McGraw-Hill, New York.

Arbib, M. A., 1972a, "The Metaphorical Brain: An Introduction to Cybernetics as Artificial Intelligence and Brain Theory," Wiley Interscience, New York.

Arbib, M. A., 1972b, Automata theory in the context of theoretical embryology, in "Foundations of Mathematical Biology," Vol. II (R. Rosen, ed.), pp. 142-215, Academic Press, New York.

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

Arbib, M. A., 1978, Segmentation, schemas and cooperative computation, in "Studies in Mathematical Biology, Part I. Cellular Behavior and the Development of Pattern" (S. Levin, ed.), Math. Assoc. of America.

Arbib, M. A., 1981a, Perceptual structures and distributed motor control, in "Handbook of Physiology, Section on Neurophysiology, Vol. III, Motor Control" (V. B. Brooks, ed.), Amer. Physiological Society, Bethesda, MD.

Arbib, M. A., 1981b, From artificial intelligence to neurolinguistics, in "Neural Models of Language Processes" (M. A. Arbib and D. Caplan, eds.).

Arbib, M. A., 1981c, Perceptual-motor processes and the neural basis of language, in "Neural Models of Language Processes" (M. A. Arbib and D. Caplan, eds.).

Arbib, M. A., and Caplan, D., 1979, Neurolinguistics must be computational, Behav. and Brain Sciences 2: 449-483.

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.), pp. 216-231, Oldenbourg.

Arbib, M. A., Lawton, D. T. and Prager, J. M. (to appear), Computing the optic flow.

Bernstein, N. A., 1968, "The Coordination and Regulation of Movements," Pergamon Press, Oxford.

Blakemore, C. and Cooper, G., 1970, Development of the brain depends on the visual environment, *Nature* 228: 477-78.

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, Computer and Information Science Department Technical Report 75C-6, 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 Behavior in the Cat, Computer and Information Science Department Technical Report 76-1, University of Massachusetts at Amherst.

Brinkman, J. and Kuypers, H. G. J. M., 1972, Split-brain monkeys: cerebral control of ipsilateral and contralateral arm, hand, and finger movements, *Science* 176: 536-539.

Brodal, A., 1969, "Neurological Anatomy," Oxford University Press.

Brodal, A., and Szikla, G., 1972, The termination of the brachium conjunctivum descendens in the nucleus reticularis tegmenti pontis: an experimental anatomical study in the cat, *Brain Res.* 39: 337-351.

Collett, T. S., 1979, A toad's devious approach to its prey: a study of some complex uses of depth vision, *J. Comp. Physiol. A*. 131: 179-189.

Cragg, B. G. and Temperley, H. N. V., 1954, The organization of neurones: a cooperative analogy, EEG *Clin. Neurophysiol*. 6: 85-92.

Craik, K. J. W., 1943, "The Nature of Explanation," Cambridge University Press.

Dev, P. 1975, Computer simulation of a dynamic visual perception model, Int. J. Man-Mach. Studies 7: 511-528.

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

Didday, R. L., 1976, A model of visuomotor mechanisms in the frog optic tectum, Math. Biosci. 30: 169-180.

Eccles, J. C., Ito, M., and Szentagothai, J., 1967, "The Cerebellum as a Neuronal Machine," Springer-Verlag.

Ede, D. A., 1978, "An Introduction to Developmental Biology," Halsted Press.

Erickson, R. P., 1974, Parallel "population" neural coding in feature extraction, in "The Neurosciences: Third Study Program" (F. O. Schmitt and F. G. Worden, eds.), The MIT Press, Cambridge, Mass.

Erman, L. and Lesser, V. R., 1980, The HEARSAY II system: a tutorial, in "Trends in Speech Recognition" (W. A. Lea, ed.), Prentice-Hall, Englewood Cliffs, NJ.

Ewert, J.-P., 1976, The visual system of the toad: behavioral and physiological studies on a pattern recognition system, in "The Amphibian Visual System: A Multidisciplinary Approach" (K. V. Fite, ed.), pp. 141-202, Academic Press, New York.

Ewert, J.-P. and von Seelen, W., 1974, Neurobiologie und System-Theorie eines visuellen Muster-Erkennungsmechanismus bei Kroten, Kybernetik 14: 167-183.

Fife, P. C., 1979, "Mathematical Aspects of Reacting and Diffusing Systems," Springer Lecture Notes in Biomathematics 28, Springer-Verlag.

Frisby, J. P. and Mayhew, J.E.W., 1977, Global processes in stereopsis, Perception 6: 195-206.

Gaze, R. M., 1970, "The Formation of Nerve Connections," Academic Press.

Gibson, J. J., 1955, The optical expansion-pattern in aerial location, Am. J. Psychol. 68: 480-484.

Gibson, J. J., 1966, "The Senses Considered as Perceptual Systems," Allen and Unwin.

Gibson, J. J., 1977, The theory of affordances, in "Perceiving, Acting and Knowing" (R. E. Shaw and J. Bransford, eds.), Erlbaum.

Grossberg, S., 1970, Neural pattern discrimination, J. Theor. Biol. 27: 291-337.

Grossberg, S., 1978, Communication, theory and development, Progr. Theor. Biol. 5: 183-232.

Grossberg, S. and Levine, D., 1975, Some developmental and attentional biases in the contrast enhancement and short term memory of recurrent neural networks, J. Theor. Biol. 53: 341-380.

Haaxma, R. and Kuypers, H., 1974, Role of occipito-frontal cortico-cortical connections in visual guidance of relatively independent hand and finger movements in rhesus monkeys, Brain Res. 71: 361-366.

Haken, H., 1978, "Synergetics: An Introduction. Nonequilibrium Phase Transitions and Self-organization in Physics, Chemistry and Biology," 2nd Edition, Springer-Verlag.

Hanson, A. R. and Riseman, E. M., 1978a, VISIONS: a computer system for interpreting scenes, in "Computer Vision Systems" (A. R. Hanson and E. M. Riseman, eds.), pp. 129-163, Academic Press, New York.

Hanson, A. R. and Riseman, E. M., eds., 1978b, "Computer Vision Systems," Academic Press, New York.

Harth, E. M., Csermely, T. J., Beek, B. and Lindsay, R. D., 1970, Brain functions and neural dynamics, J. Theor. Biol. 26: 93-120.

Herman, R. M., Grillner, S., Stein, P.S.G. and Stuart, D. G., eds., 1976, "Neural Control of Locomotion," Plenum Press.

Hirsch, H.V.B. and Spinelli, D. N., 1971, Modification of the distribution of receptive field orientation in cats by selective visual exposure during development, Exp. Brain Res. 12: 509-527.

Hope, R. A., B. J. Hammond and Gaze, R. M., 1976, The arrow-model of retino-tectal connections, Proc. Roy. Soc. Lond. Ser. B 194: 447-466.

Horn, B.K.P., 1974, Determining lightness from an image, Computer Graphics and Image Understanding 3: 277-299.

Hubel, D. H. and Wiesel, T. N., 1962, Receptive fields, binocular interaction and functional architecture in the cat's visual cortex, J. Physiol. 160: 106-154.

Hubel, D. H. and Wiesel, T. N., 1965, Binocular interaction in striate cortex of kittens reared with artificial squint, J.

Neurophysiol. 28: 1041-1059.

Hubel, D. H. and Wiesel, T. N., 1974, Sequence regularity and geometry of orientation columns in the monkey striate cortex, *J. Comparative Neurology* 158: 267-274.

Humphrey, N. K., 1970, What the frog's eye tells the monkey's brain, *in* "Subcortical Visual Systems" (D. Ingle and G. E. Schneider, eds.), pp. 324-337, Karger.

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

Ingle, D., 1976, Spatial vision in anurans, *in* "The Amphibian Visual System" (K. V. Fite, ed.), pp. 119-141, Academic Press, New York.

Jackson, J. H., 1874, On the nature of the duality of the brain, *Med. Press and Circular* 1: 19-63.

Jackson, J. H., 1878-79, On affections of speech from disease of the brain, *Brain* 1: 304-330, 2: 203[illegible], 323-356.

Jeannerod, M. and Biguer, B., 1980, Visuomotor mechanisms in reaching within extra-personal space, *in* "Advances in the Analysis of Visual Behavior" (D. J. Ingle, R. J. W. Mansfield and M. A. Goodale, eds.), The MIT Press, Cambridge, Mass.

Julesz, B., 1971, "Foundations of Cyclopean Perception," University of Chicago Press.

Katchalsky, A., Rowland, V. and R. Blumenthal, 1974, "Dynamic Patterns of Brain Cell Assemblies," The MIT Press, Cambridge, Mass.

Kilmer, W. L., McCulloch, W. S., and Blum, J., 1969, A model of the vertebrate central command system, *Int. J. Man-Machine Stud*. 1: 279-309.

Kohonen, T. and Oja, E., 1976, Fast adaptive formation of orthogonalizing filters and associative memory in recurrent networks of neuron-like elements, *Biol. Cybernet*. 21: 85-95.

N. Kopell, 1978, Reaction-diffusion equations and pattern formation, *in* "Studies in Mathematical Biology, Part I. Cellular Behavior and the Development of Pattern" (S. A. Levin, ed.), pp. 191-205, Math. Assoc. of America.

Kuffler, S. W., 1953, Discharge patterns and functional organization of mammalian retina, J. *Neurophysiol*. 16: 37-68.

Lara, R., 1980, The amphibian visual system: modelling the tectal column, *Cognition and Brain Theory* 3: 70-100.

Lara, R. Arbib, M. A., and Cromarty, A. S. (to appear), The role of the tectal column in facilitation of amphibian prey-catching behavior: a neural model.

Lavorel, P. M. and Arbib, M. A. (to appear), Towards a theory of language performance: neurolinguistics, perceptual-motor processes, and cooperative computation, Theoretical Linguistics.

Lee, D. N., 1974, Visual information during locomotion, in "Perception: Essays in Honor of James J. Gibson" (R. B. MacLeod and H. L. Pick, Jr., eds.), pp. 250-267, Cornell Unviersity Press.

Lee, D. N. and Lishman, J. R., 1977, Visual control of locomotion, Scand. J. Psychol. 18: 224-230.

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

Levin, S. A. (ed.), 1978, "Studies in Mathematical Biology, Part I. Cellular Behavior and the Development of Pattern," Math. Assoc. of America.

Lock, A. and Collett, T., 1980, The three-dimensional world of a toad, Proc. Roy. Soc. Lond. B206: 481-487.

MacKay, D. M., 1966, Cerebral organization and the conscious control of action, in "Brain and Conscious Experience" (J. C. Eccles, ed.), pp. 422-440, Springer-Verlag.

Marr, D., 1978, Representing visual information, in "Computer Vision Systems" (A. R. Hanson and E. M. Riseman, eds.), pp. 61-80, Academic Press, New York.

Marr, D. and Poggio, T., 1977a, From understanding computation to understanding neural circuitry, Neurosci. Res. Prog. Bull. 15: 470-488.

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

Marr, D. M. and Poggio, T., 1979, A theory of human stereopsis, Proc. Roy. Soc. Ser. B 204: 301-328.

Marr, D., Palm, G. and Poggio, T., 1978, Analysis of a cooperative stereo algorithm, Biol. Cybernetics 28: 223-239.

Minsky, M. L., 1975, A framework for representing knowledge, in "The Psychology of Computer Vision" (P. H. Winston, ed.), pp. 211-277, McGraw-Hill.

Mountcastle, V. B., 1957, Modality and topographic properties of single neurons of cat's somatic sensory cortex, J. Neurophysiol. 20: 408-434.

Mountcastle, V. B., 1978, An organizing principle for cerebral function: the unit module and the distributed system, in G. M. Edelman and V. B. Mountcastle, "The Mindful Brain," The MIT Press, Cambridge, Mass.

Murphy, J. T. and Sabah, N. H., 1970, The inhibitory effect of climbing fiber activation on cerebellar Purkinje cells, Brain Res. 19: 486-490.

Neisser, U., 1976, "Cognition and Reality," W. H. Freeman.

Orlovsky, G. N., 1972, The effect of different descending systems on flexor and extensor activity during locomotion, Brain Res. 40: 359-372.

Oscarrson, O., 1973, Functional organization of spinocerebellar paths, in "Handbook of Sensory Physiology, Vol. II. Somatosensory System" (A. Iggo, ed.), pp. 339-380, Springer-Verlag.

Overton, K. J. and Arbib, M. A. (to appear), The XBAM model of retino-tectal connections.

Parma, C. C., Hanson, A. R. and Riseman, E. M., 1980, Experiments in schema-driven interpretation of a natural scene, Department of Computer and Information Science Technical Report 80-10, University of Massachusetts at Amherst.

Perkel, D. H. and Bullock, T. H., 1968, "Neural Coding," Neurosci. Res. Program Bulletin 6(3).

Prager, J. M., 1979, Segmentation of Static and Dynamic Scenes, Ph.D. Thesis, Computer and Information Science Department, University of Massachusetts at Amherst.

Prager, J. M. and Arbib, M. A. (to appear), Computing the optic flow: the MATCH algorithm and prediction.

Rashevsky, N., 1948, "Mathematical Biophysics," 2nd edition, University of Chicago Press.

Rinzel, J., 1978, Integration and propagation of neuroelectric signals, in "Studies in Mathematical Biology, Part I. Cellular Behavior and the Development of Pattern" (S. A. Levin, ed.), pp. 1-66, Math. Assoc. of America.

Scheibel, M. E. and Scheibel, A. B., 1958, Structural substrates for integrative patterns in the brain stem reticular core, in "Reticular Formation of the Brain" (H. H. Jasper et al., eds.), pp. 31-68, Little, Brown and Co.

Shaw, G. L., 1978, Space-time correlations of neuronal firing related to memory storage capacity, Brain Res. Bull. 3: 107-113.

Sherrington, C. S., 1910, Flexion-reflex of the limb, crossed extension-reflex, and reflex stepping and standing, J. Physiol. 40: 28-121.

Shik, M. L., Severin, F. V. and Orlovskii, G. N., 1966, Control of walking and running by means of electrical stimulation of the mid-brain, Biophysics 11: 756-765.

Sperling, G., 1970, Binocular vision: a physical and neural theory, Am. J. Psych. 83: 463-534.

Sperry, R. W., 1944, Optic nerve regeneration with return of vision in anurans, J. Neurophysiol. 7: 57-70.

Spinelli, D. N. and Jensen, F. E., 1979, Plasticity: the mirror of experience, Science 203: 75-78.

Szentagothai, J., 1978, The neuron network of the cerebral cortex: a functional interpretation, Proc. Roy. Soc. Lond. B201: 219-248.

Szentagothai, J. and Arbib, M. A., 1975, "Conceptual Models of Neural Organization," The MIT Press, Cambridge, Mass.

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

Turing, A. M., 1952, The chemical basis of morphogenesis, Phil. Trans. Roy. Soc. Lond. 237: 37-72.

Ullman, S., 1979, "The Interpretation of Visual Motion," The MIT Press, Cambridge, Mass.

von der Malsburg, C., 1973, Self-organization of orientation-sensitive cells in the striate cortex, Kybernetik 14: 85-100.

von der Malsburg, C., and Willshaw, D. J., 1977, How to label nerve cells so that they can interconnect in an orderly fashion, Proc. Natl. Acad. Sci. USA 74: 5176-5178.

von Neumann, J., 1951, The general and logical theory of automata, in "Cerebral Mechanisms in Behavior: The Hixon Symposium" (L. A. Jeffress, ed.), pp. 1-32, Wiley.

Voogdt, J., 1969, The importance of fiber connections in the comparative anatomy of the mammalian cerebellum, in "Neurobiology of Cerebellar Evolution and Development" (R.

Llinas, ed.), pp. 493-514, American Medical Assoc.

Wiener, N., 1948, "Cybernetics: or Control and Communication in the Animal and the Machine," 1st edition, Wiley.

Wilson, W. R. and Cowan, J. D., 1973, A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue, Kybernetik 13: 55-80.

York, B. W., 1980, Shape Representation in Computer Vision, Ph.D. Thesis, Computer and Information Science Department, University of Massachusetts at Amherst.

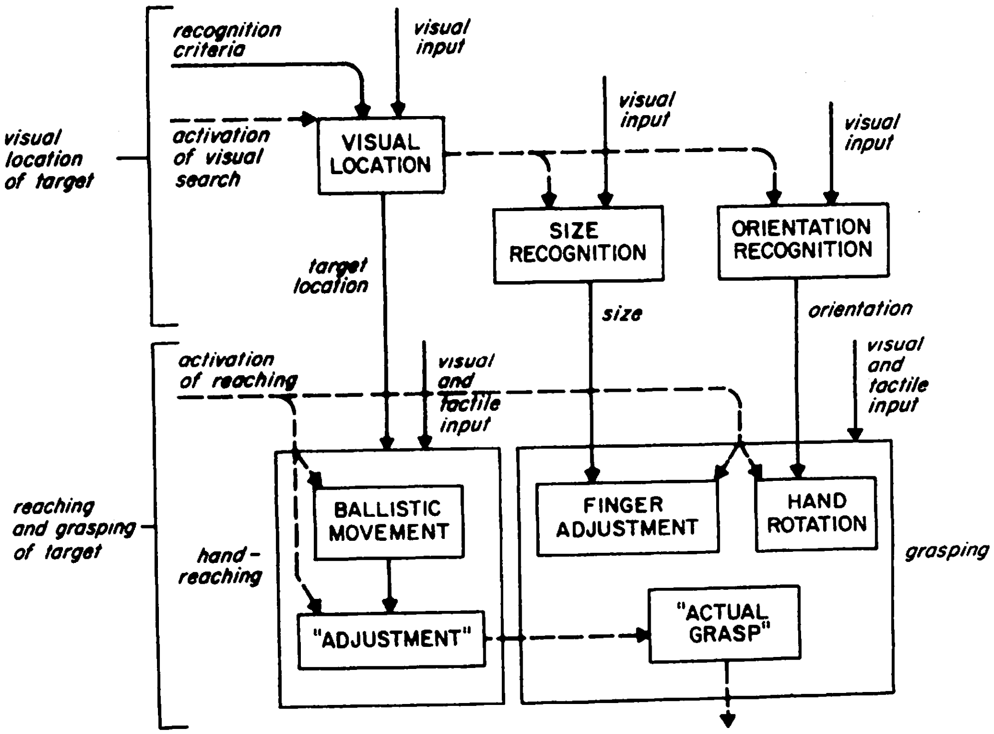

Figure 1. A coordinated control program for grasping an object.

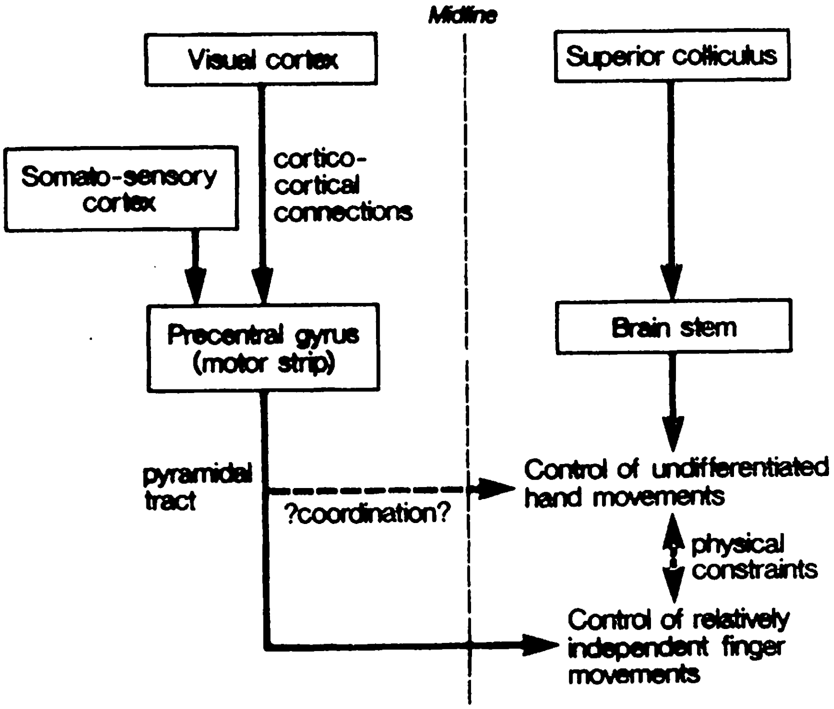

Figure 2. Pyramidal pathway 'differentiates' extrapyramidal control of reaching movements.

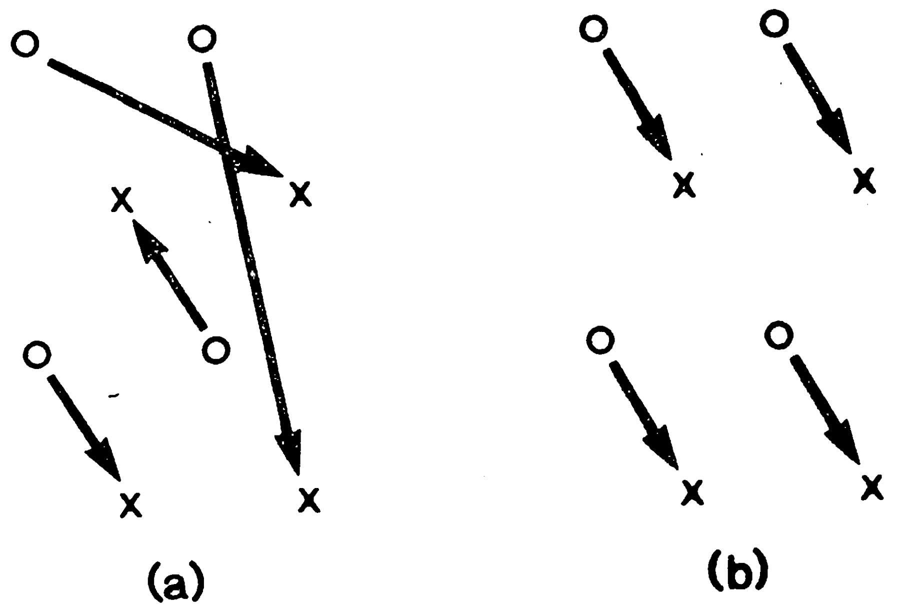

Figure 3. In a world made up of surfaces, nearby features are likely to have similar optic flow. Thus the flow of (b) is far more likely to be correct than that of (a).

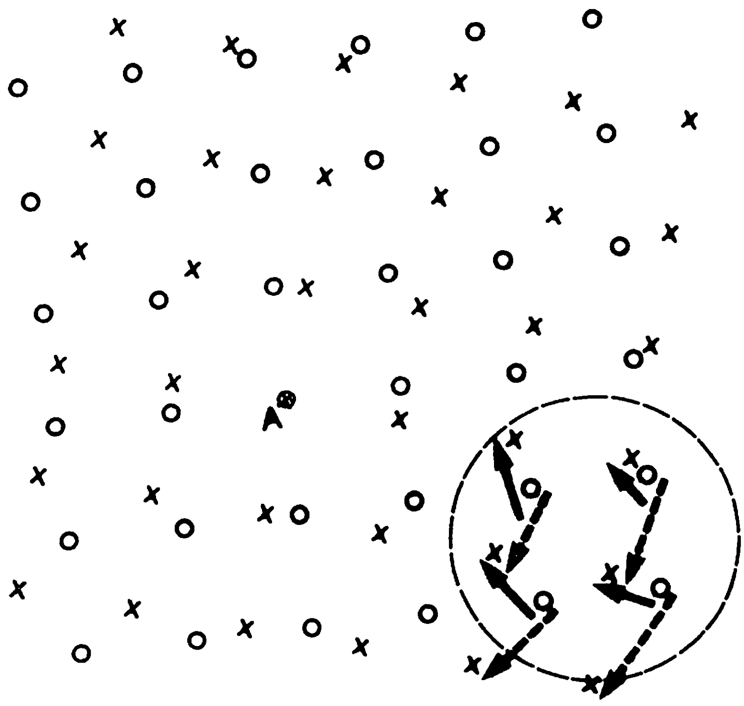

Figure 4. Frame 1 comprises the dots indicated by circles; Frame 2 is obtained by rotating the array about the pivot at A to place the dots in the positions indicated by crosses. The dashed circle at lower right is the receptive field of a local processor. The solid arrows indicate the best local estimate of the optic flow, the dashed arrows show that actual pairing of features under rotation about A.

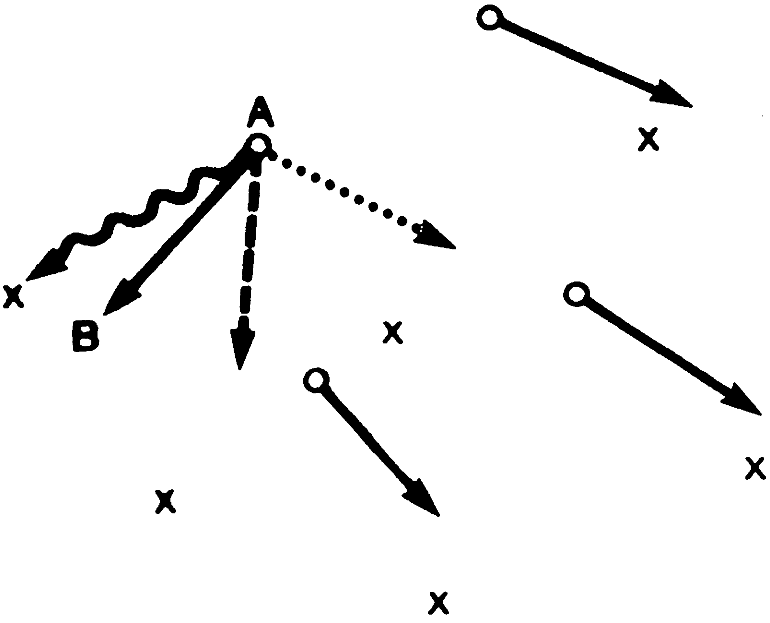

Figure 5. The circles indicate features in Frame 1, the crosses features in Frame 2, and the solid arrows the current estimate of the optic flow -- the head of the arrow shows the posited position in Frame 2 of the feature corresponding to the Frame 1 feature at the tail of the arrow. 'Feature matching' alone would adjust A's optic flow to the wavy arrow pointing to the Frame 2 feature nearest to B (the current estimate of A's Frame 2 position); 'local smoothness' would yield the dotted arrow, the average of the optic flow of the neighbors; while our relaxation algorithm yields the dashed arrow as a weighted combination of these two estimates.

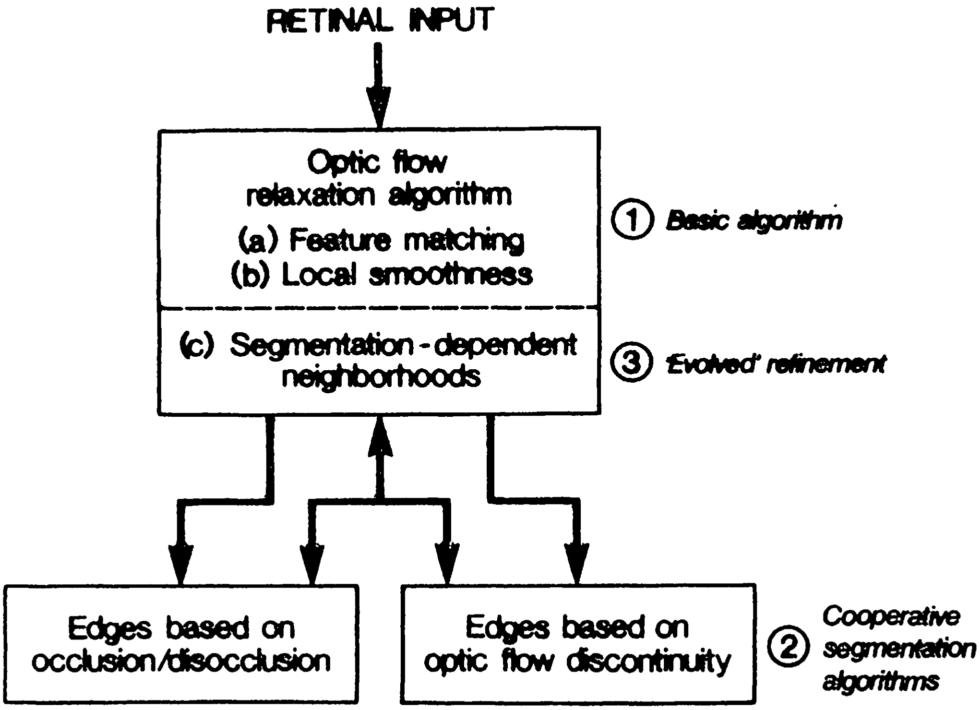

Figure 6. (1) Our basic optic flow relaxation algorithm uses the consistency conditions of feature matching and local smoothing. (2) The resultant optic flow estimate permits the hypothesization of edges on cues based on both occlusion/disocclusion cues and on optic flow discontinuity. (3) The resultant edges hypotheses can be used to refine the computation of optic flow by dynamically adjusting the neighborhoods used in employing the consistency conditions.

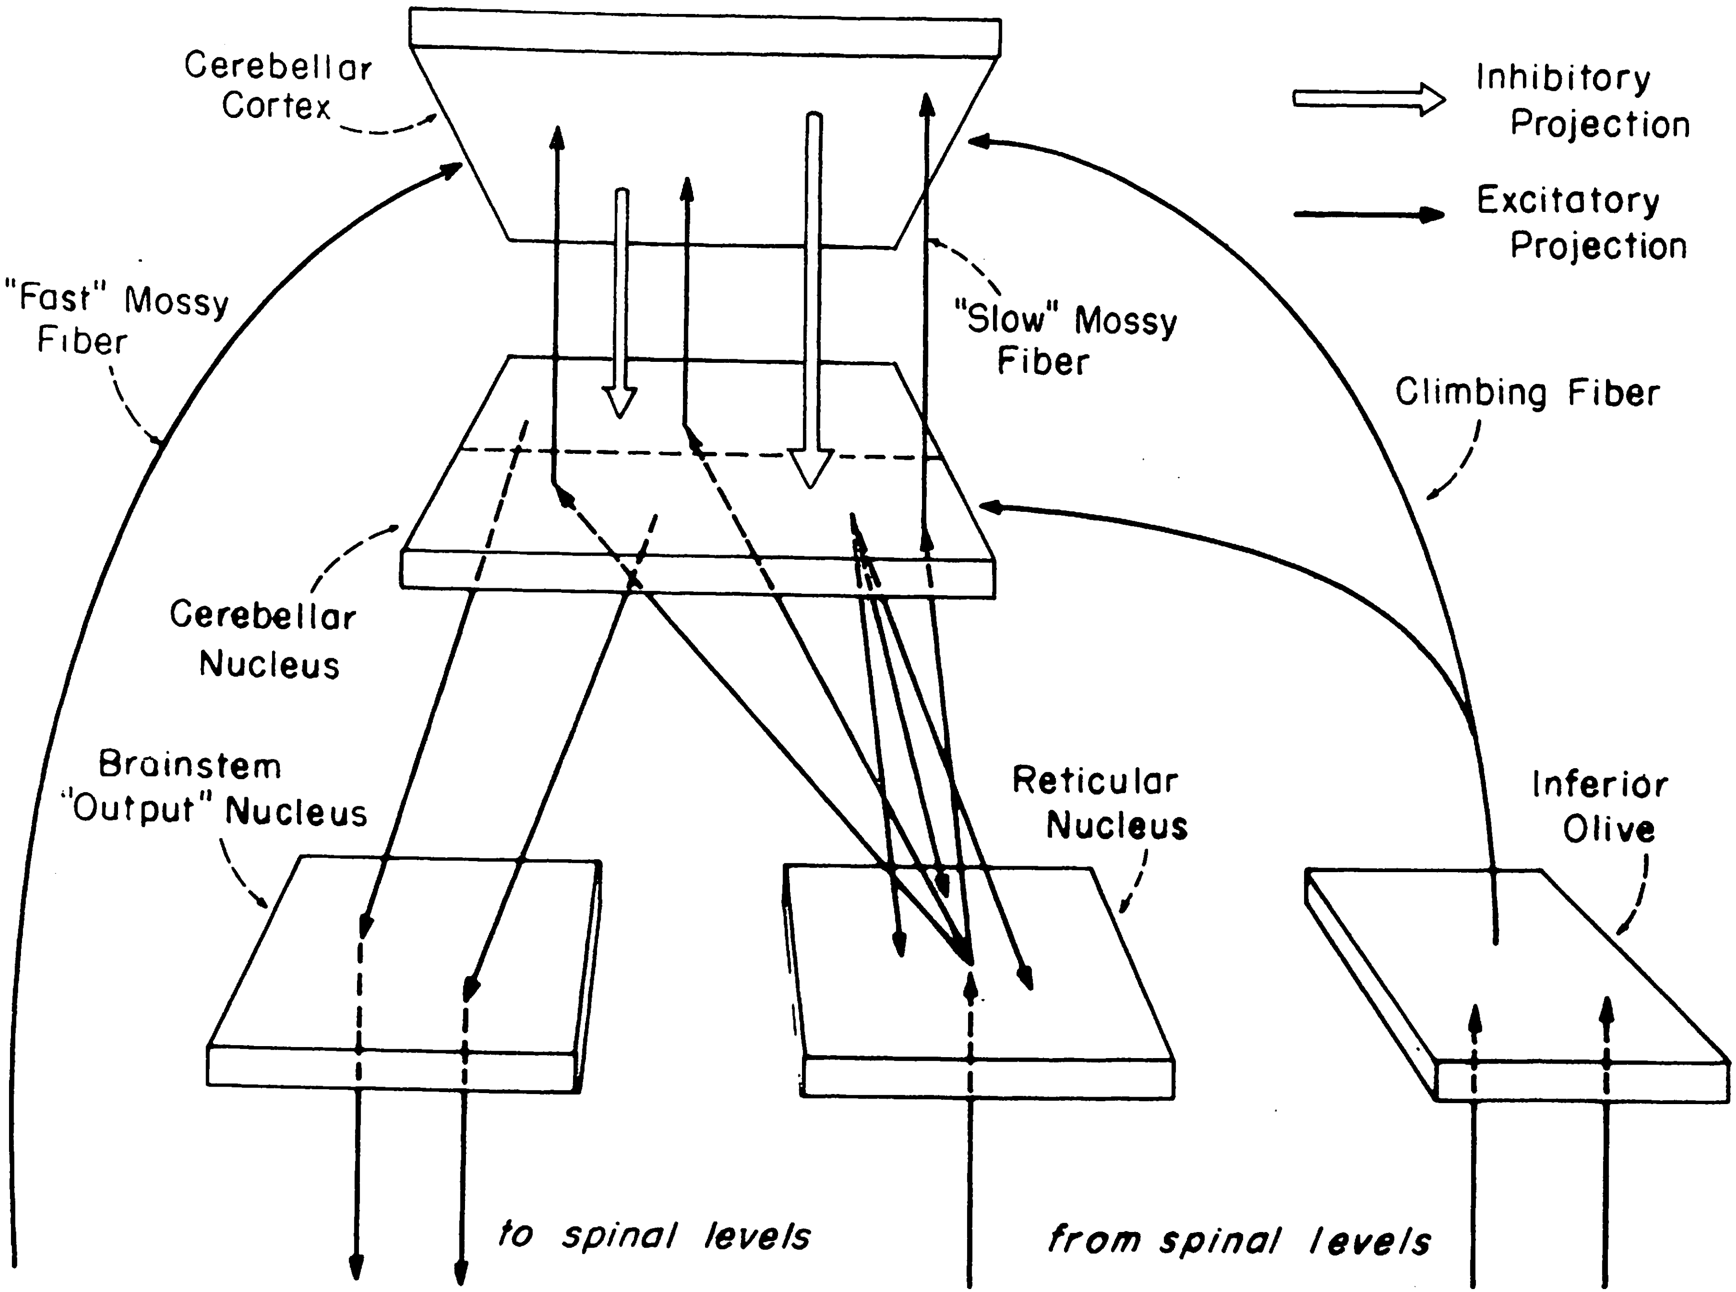

Figure 7. Schematic of the interacting control surfaces involved in tuning of motor schemas by the cerebellum and related nuclei.

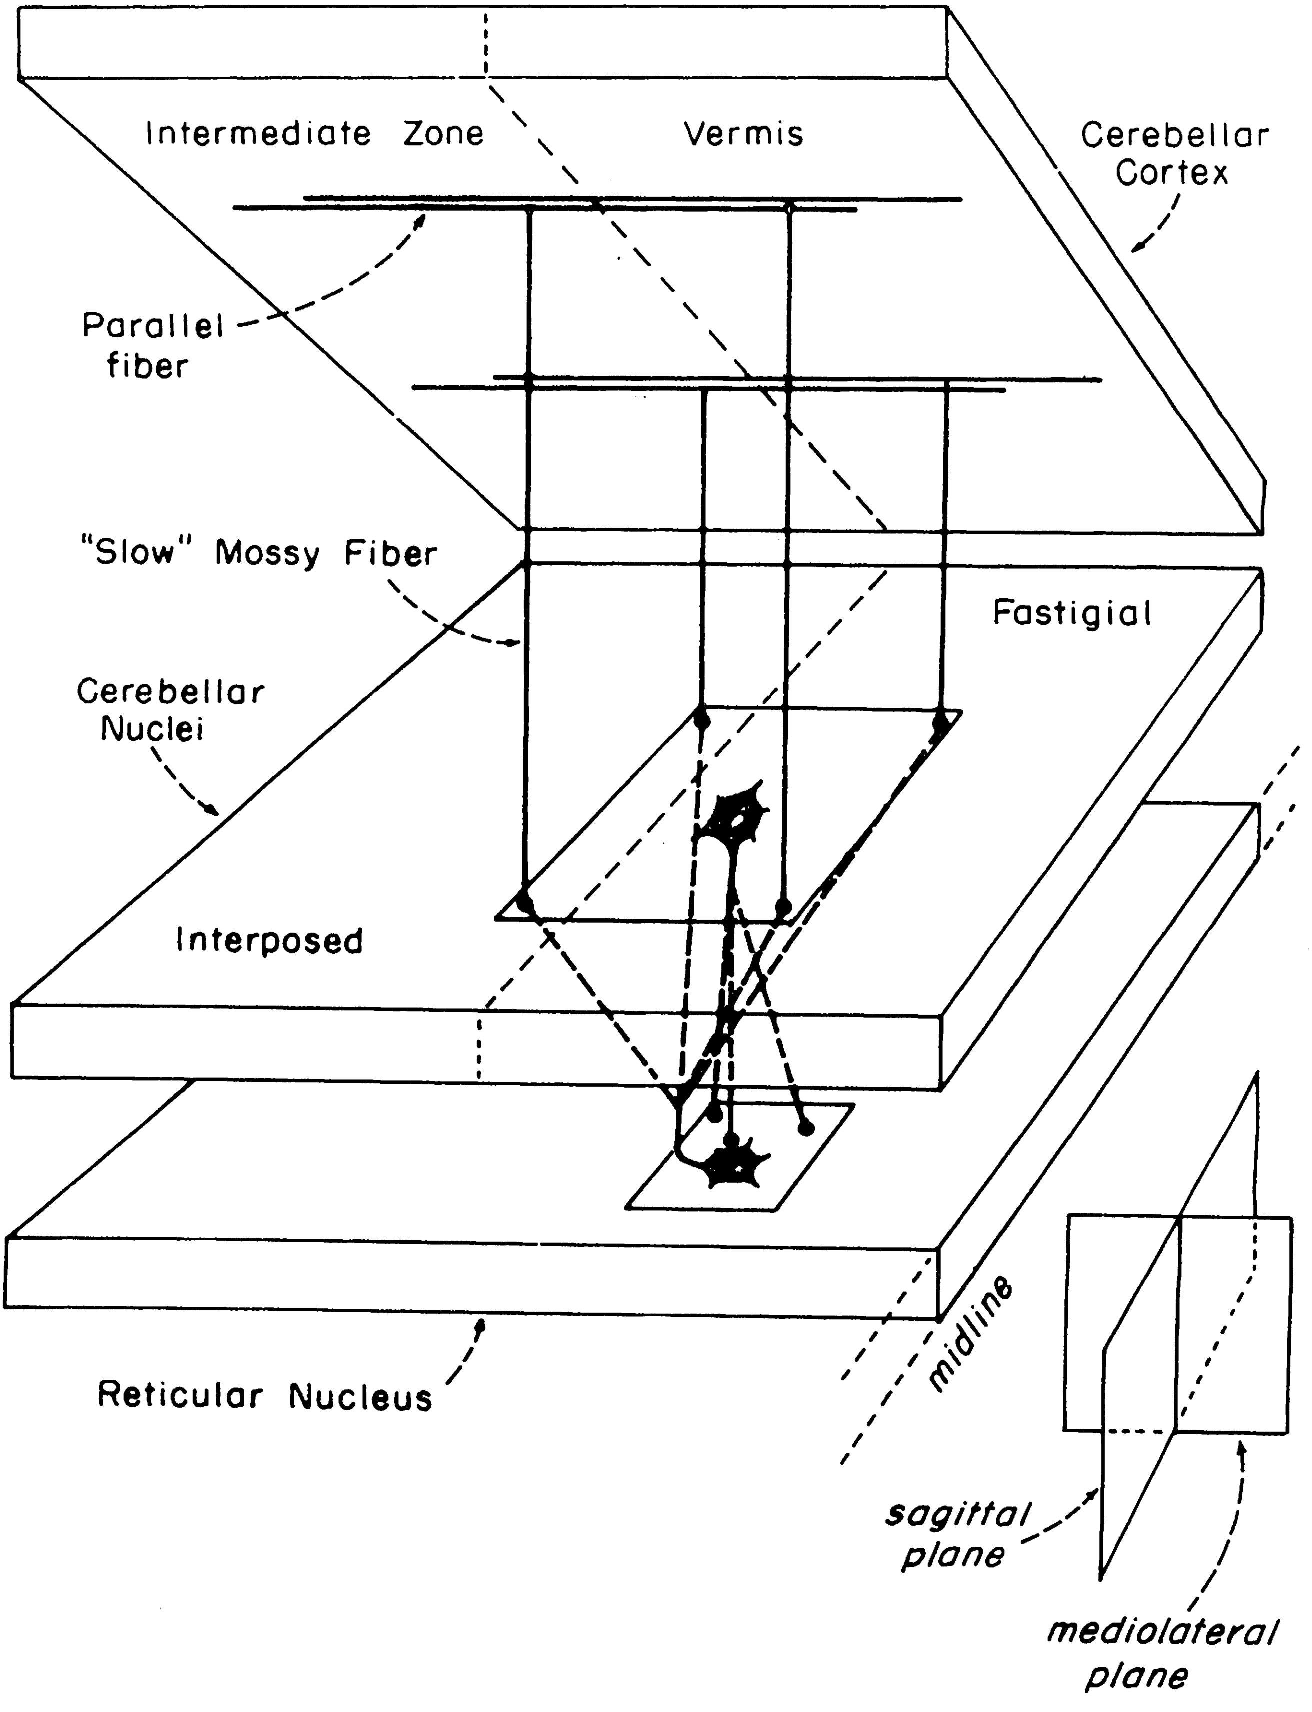

Figure 8. Anatomical template of cerebello-reticular reverberatory loop.

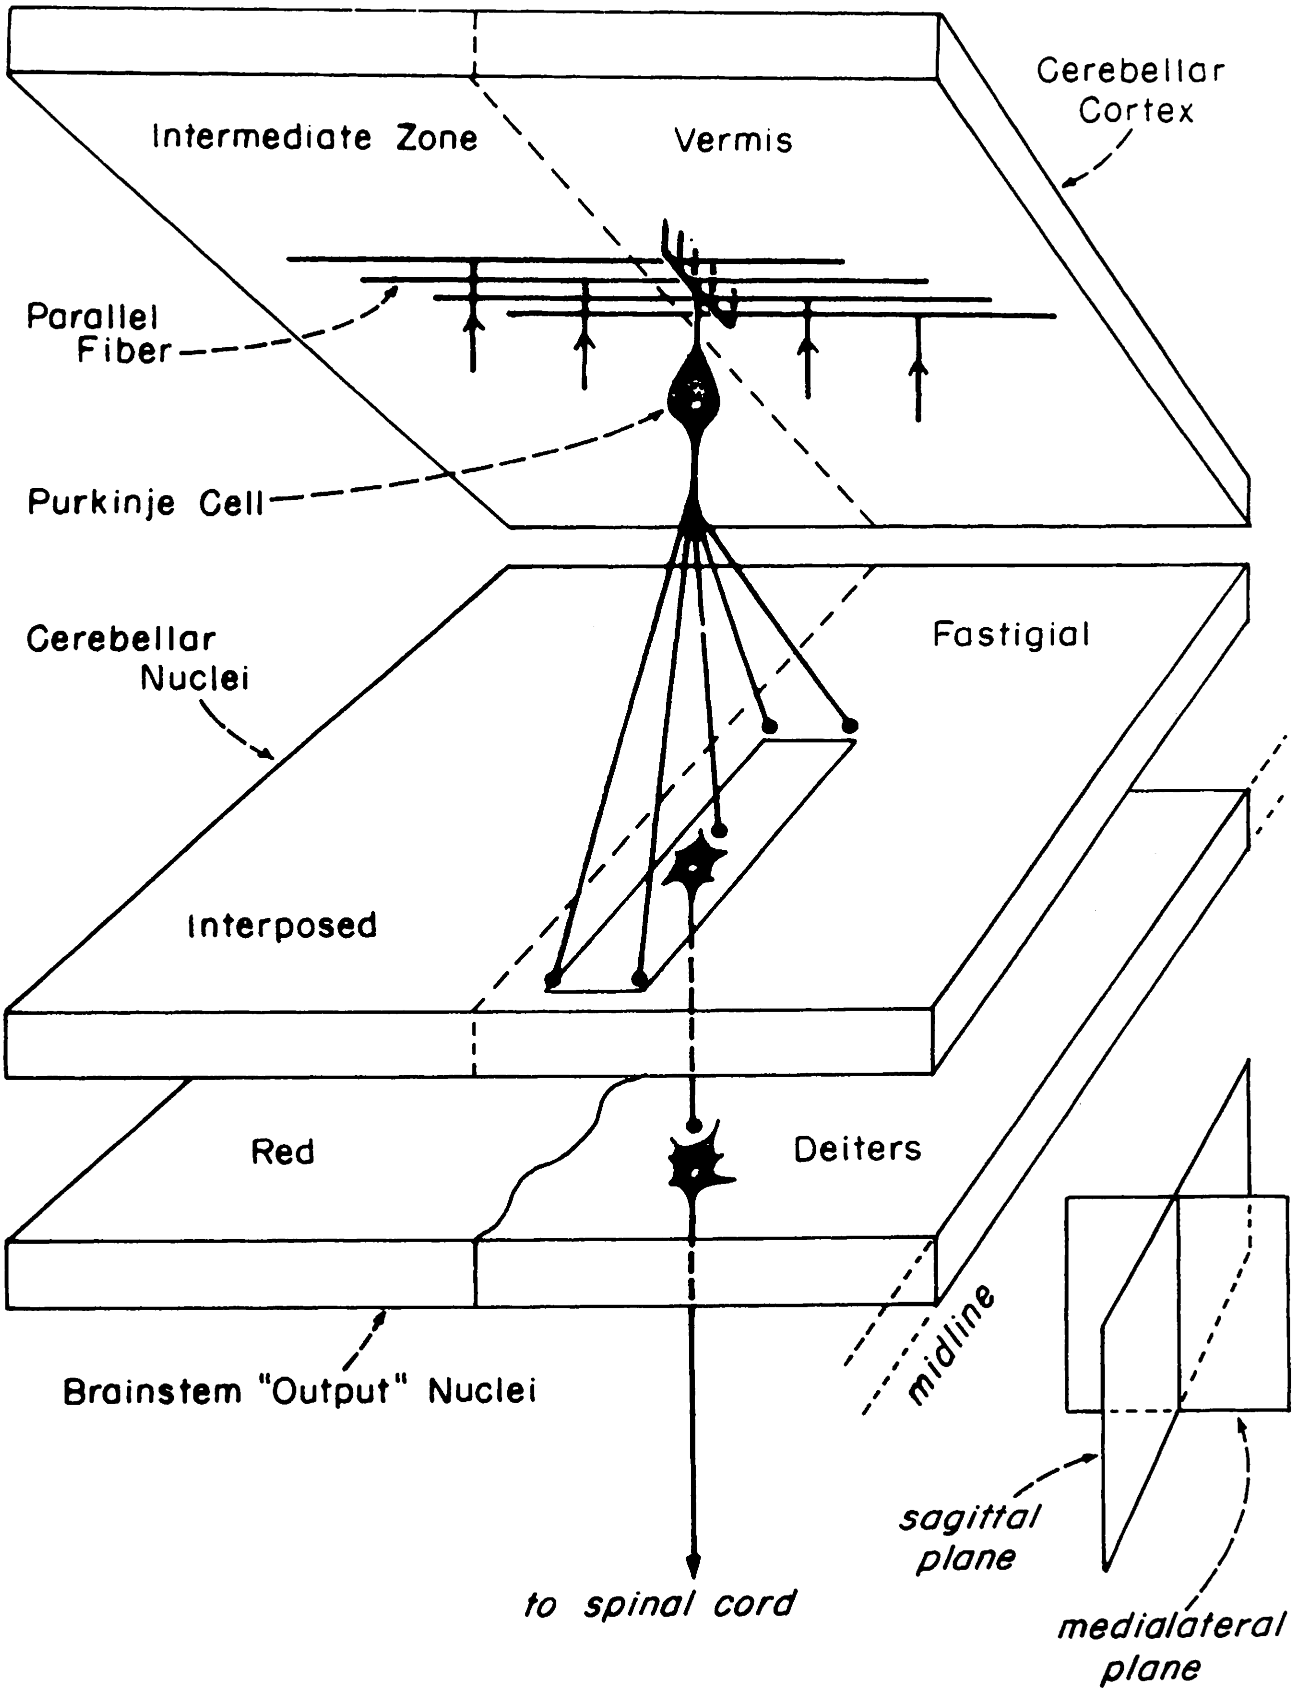

Figure 9. Template of cortico-nuclear projection and cerebellar outflow.

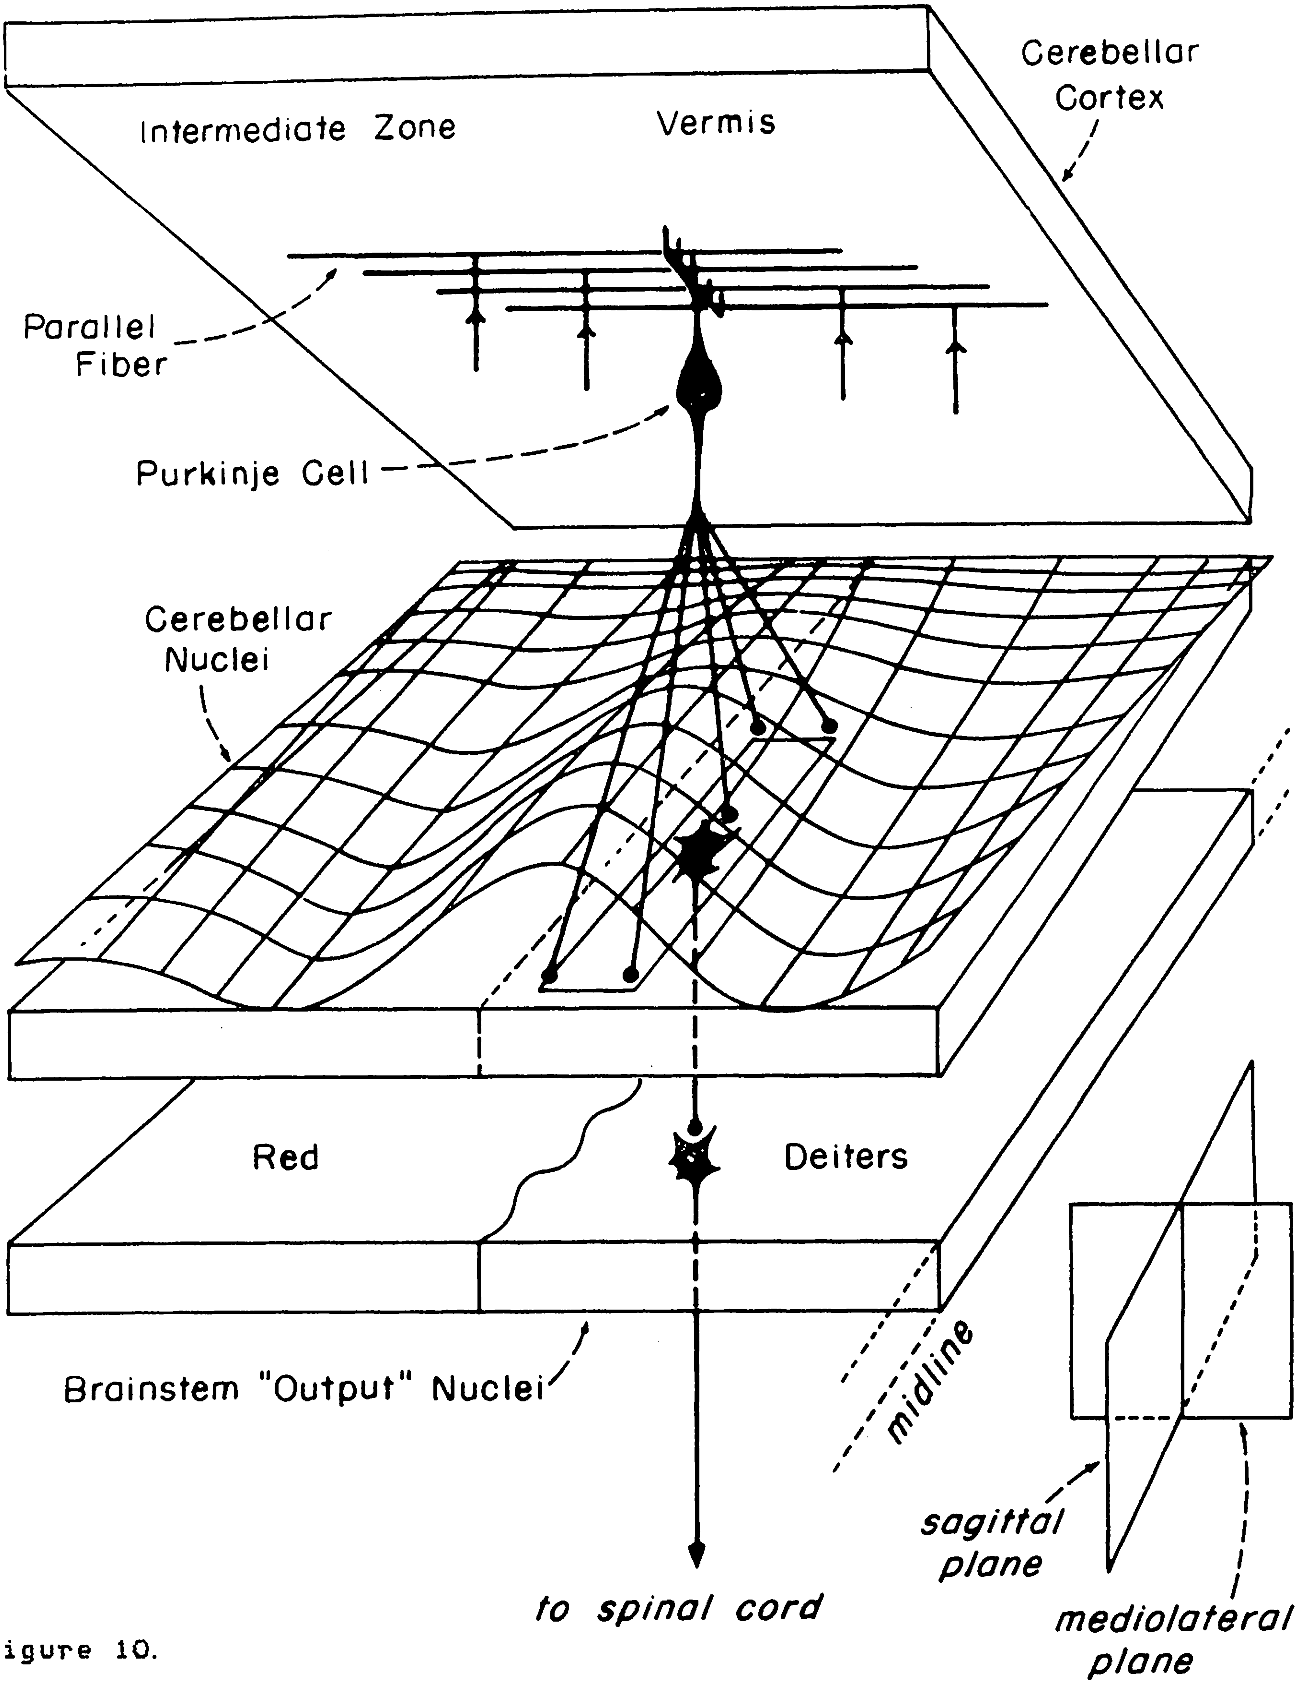

Figure 10.

Activation pattern 'stored' in cerebellar neuronal interactions via climbing fiber activity represents the parameters for a particular motor schema which may activate the musculature.

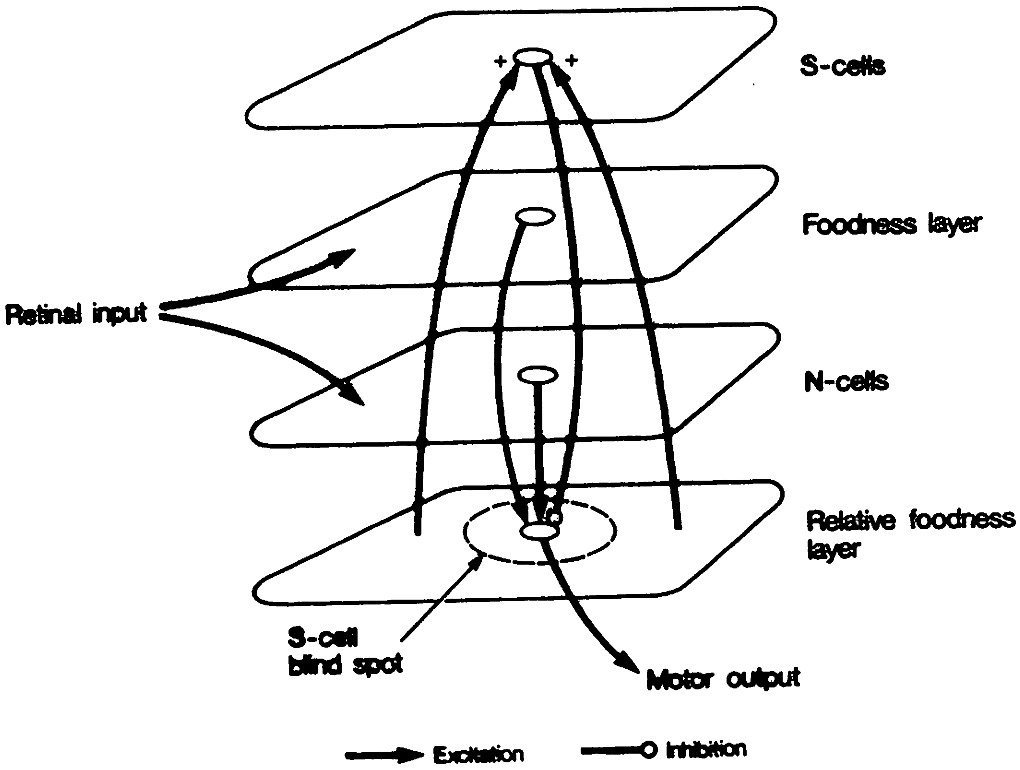

Figure 11. Tectal-pretectal interaction in model of prey-selection in frog.