

PERCEPTUAL STRUCTURES AND
DISTRIBUTED MOTOR CONTROL

Michael A. Arbib
Computer and Information Science, and
Center for Systems Neuroscience
University of Massachusetts
Amherst, MA 01003

COINS Technical Report ⁷⁹~~78~~-11

(June 1979)

This report is a preprint of a chapter for

Volume III on Motor Control,
edited by Vernon B. Brooks,

of the

Handbook of Physiology

to be published in 1980 by

The American Physiological Society
9650 Rockville Pike
Bethesda, Maryland 20014

CHAPTER CONTENTS

PERCEPTUAL STRUCTURES AND
DISTRIBUTED MOTOR CONTROL

Michael A. Arbib
 Computer and Information Science, and
 Center for Systems Neuroscience
 University of Massachusetts
 Amherst, MA 01003

COINS Technical Report ⁷⁹~~80~~11

(June 1979)

This report is a preprint of a chapter for

Volume III on Motor Control,
 edited by Vernon B. Brooks,
 of the

Handbook of Physiology

to be published in 1980 by

The American Physiological Society
 9650 Rockville Pike
 Bethesda, Maryland 20014

Introduction

The Scope of the Chapter

The Place of Brain Theory within Cybernetics

Concepts from Computer Science and Control Theory

Programs need not be stereotypes

State and feedback in control theory

Serial order in behavior revisited

Visuomotor Coordination in Frog and Toad

Maps as Control Surfaces

A distributed processing model of superior colliculus

Cooperative computation in somatotopically organized networks

Distributed motor control

A Model of Frog Snapping

A tectal model of choice behavior

Competition and cooperation in neural nets

The Many Visual Systems

Cooperative computation of controllers

Prey-enemy pattern recognition

Summary

Perceptual and Motor Schemas

Perceptual Schemas and the Action/Perception Cycle

The internal model

Affordances

The action-perception cycle

Planning
 Perception as potential action
 Coordinative structures
 The schema-assemblage
 Eye-movements and visual perception
 Optic Flow and the Control of Movement
 Focus of expansion
 Relative motion
 Inferring collision information
 Time until adjacency
 Motor Schemas
 Synergies
 System adaptation
 Mutability of motor schemas
 Summary

 Coordinated Control Programs
 Feedforward
 Discrete-activation feedforward
 Co-activation feedforward
 Motor schemas again
 Interwoven Activation of Motor Schemas
 A program for grasping
 Feedforward and the cerebellum
 When is a program "hard-wired"?
 Competition and cooperation of motor schemas

Skill Acquisition
 Summary

 The Perspective of Artificial Intelligence
 Programs and Planners
 Compilers and interpreters
 Planning techniques in artificial intelligence
 The diversity of programs
 Mapping the environment
 The monkey and the banana
 Maps and topologies
 Developing a high-level vocabulary
 Program Synthesis and Visuomotor Coordination
 An experiment on visuomotor coordination
 Program synthesis and the available repertoire
 From set to topological space
 Summary

 Conclusion

Footnote of Acknowledgement

The preparation of this chapter was supported in part by NIH under grant number 1 R01 NS14971-01. I am most grateful to many friends and colleagues for their comments on two earlier drafts. Curt Boylls and Richard Reiss saved me from many indiscretions. Gwyn Mitchell did a fine job of typing. Finally, my thanks to my family for helping me get through to the light at the end of the tunnel.

INTRODUCTION

The Scope of the Chapter

The overall goal of this chapter is to offer a conceptual analysis of the internal structures and processes active in an organism using visual input to guide its interaction with a complex environment. As such, the chapter will range from structures and processes well rooted in the experimental laboratory to those which appear logically necessary for the control of complex behavior but which have so far proved resistant to experimental analysis. However, even in these latter cases, various lines of rapprochement between theory and experiment will be briefly sketched.

The chapter, then, is designed to complement the "lumped model" approach to biological control systems analysis of the α - γ system, the control of eye movements, and so on. Studies of this kind have made biological control theory the focus for successful integration of mathematical analysis with both behavioral and neural experimentation in seeking neural mechanisms for sensorimotor coordination. But we can see that the usual approaches are limited in several ways:

(a) Sensory input is usually lumped into a few variables such as the angular orientation of a target for visual tracking. We need an analysis of rich perceptual structures which represent a structured environment in a fashion compatible with the behavioral repertoire of the organism.

(b) The emphasis on lumped control systems has downgraded the analysis of such systems as played out over layered neural structures. We need to analyze how maps in the brain may be viewed as control surfaces for such layered control systems.

(c) The study of specific skills or reflexive behaviors has drawn attention from the fact that much behavior (even animal, but especially human) is not the product of a single innate control system. Rather, we need to analyze the planning of behavior in terms of coordinated control programs.

The above list in no way invalidates study of single lumped control systems, since the choice of appropriate simplifications is the key to good theory and experiment. But the use of a single target variable may, for example, miss the essence of tectal-pretectal interactions in the frog if a task of the system is to guide the frog's action in a complex structured environment which contains several flies and predators. And features of a brain region which seem inexplicable when it is viewed as an isolated control structure may be clarified when we realize that control is distributed over many such structures, and much of the architecture of each region is dictated by the need for proper coordination. The section on visuomotor coordination in frog and toad will address many of these issues at a level which makes rich contacts with neuroscientific experiments.

In the section on perceptual and motor schemas, we briefly examine the way in which the animal's "internal representation of the world" may be viewed as an assemblage of schemas (each corresponding to some object or domain of interaction in the environment) which attunes the animal for possible interactions with its environment. This representation is in a continual state of flux as the cycle of action and perception is a constant of the animal's behavior. Rather than study perception in general, we discuss the way in which a time-changing pattern of visual input (as distinct from a localized trigger feature) can be analyzed, with optic flow yielding

parameters necessary for the control of movement. We then study motor schemas -- akin to the synergies of the Russian school -- as the "units" of motor control, turning to the experimental literature for evidence that such "units" do exist, but are by no means immutable.

In the section on coordinated control programs, we first re-examine feedforward, finding it useful to distinguish discrete-activation and co-activation. We view a motor schema as being a system combining feedback and feedforward with an "identification procedure" which can adapt its parameters to a changing environment. We then examine hypothetical programs for coordinating the activation of such motor schemas.

It is clear that, in human behavior at least, many programs for motor control are not "hard-wired" but are rather the result of planning attuned to current circumstances. The study of such planning processes has so far yielded little to experiment. To catalyze the development of an appropriate vocabulary we devote a section to perspectives from that field of computer science known as artificial intelligence, studying a number of planning programs, and offering a hypothetical account of the role program synthesis might play in the acquisition of visuomotor coordination.

The theoretical analysis of perceptual structures and distributed motor control is still in its infancy, while the experimental analysis of such models lags even further behind. It would be premature, then, to insist that the concepts developed here must become part of the neuroscientist's repertoire. But it does seem proper to insist that the questions raised are important for neuroscience, and that formal tools -- both theoretical and experimental -- are badly needed in their elucidation. It is in this spirit of encouraging the mutual exploration of these questions by theorist and experimentalist that the concepts of this chapter are offered.

The Place of Brain Theory within Cybernetics

Technology has always played a crucial role in attempts to understand the human mind and body: the study of the steam-engine contributed concepts for the study of metabolism, and electricity has been part of the study of the brain at least since Galvani touched frog leg to iron railing sometime before 1791. In 1748, Julien La Mettrie published L'Homme machine (88) and suggested that such automata as the mechanical duck and flute player of Vaucanson indicated the possibility of one day building a mechanical man that could talk. The automata of that day were unable to adapt to changing circumstances, but in the following century machines were built which could automatically counter disturbances to restore the desired performance of the machine -- perhaps the best known example being Watt's governor for his steam engine. This led to Maxwell's 1868 paper "On Governors" (78) which laid the basis for what we would now call the theory of negative feedback, as well as the study of system stability. At the same time, Claude Bernard (16) was drawing attention to what, sixty years later, Walter Cannon (24) would dub homeostasis, when he observed that physiological processes often form circular chains of cause and effect which could serve to counteract disturbances in such variables as body temperature, blood pressure, and glucose level in the blood (see (13) for the relation of homeostasis to cybernetics).

1943 was the key year for bringing together the notions of control mechanism and intelligent automata. K.J.W. Craik, whose posthumously published papers (27, 28, 29, 30) were to have such an influence on the control-theoretic analysis of human skills, published his seminal essay on "The Nature of Explanation" (26) in which the nervous system was viewed "as a calculating machine capable of modelling or paralleling external events",

suggesting that this process of paralleling is the basic feature of thought and explanation. In the same year, Rosenblueth, Wiener and Bigelow (130) published "Behavior, Purpose and Teleology." Engineers had noted that if feedback used in controlling the rudder of a ship, say, is too brusque, the rudder will overshoot, compensatory feedback will yield a larger overshoot in the opposite direction, and so again and again as the system goes into wild oscillations. Wiener and Bigelow had asked Rosenblueth if these were any corresponding pathological condition in humans, and he had offered the example of intention tremor associated with injury to the cerebellum. This evidence for feedback within the human nervous system led the three scientists to urge that neurophysiology should move beyond the Sherringtonian view of the central nervous system as a reflex device adjusting itself in response to sensory inputs. Rather, the concept of setting of reference values for feedback systems could provide the basis for the analysis of the brain as a purposive system explicable only in terms of circular processes, from nervous system to muscles to the external world returning thence by receptors, whether they be exteroceptors or proprioceptors.

1943 also saw the publication of "A Logical Calculus of the Ideas Immanent in Nervous Activity" by McCulloch and Pitts (99) in which they offered their formal model of the neuron as a threshold logic unit, building on the neuron doctrine of Ramon y Cajal, and the excitatory and inhibitory synapses of Sherrington. They used notation from the mathematical logic of Whitehead and Russell and Carnap, but a major stimulus for their work was the Turing machine, a device which could read, write and move upon an indefinitely extendible tape each square of which bore a symbol from some finite alphabet. Turing (142) had made plausible the claim that any effectively definable computation (i.e. anything that a human could do in the way of symbolic manipulation by following a finite and completely explicit set of

rules -- in 1936 the referent of 'computer' was still a human!) could be carried out by such a machine equipped with a suitable program. What McCulloch and Pitts demonstrated was that each such program could be implemented using a finite network, with loops, of their formal neurons. Thus, as electronic computers came to be built toward the end of World War II, it was understood that whatever they could do could be done by a network of neurons.

These, then, were some of the strands which were gathered in Wiener's 1948 book "Cybernetics: or Control and Communication in the Animal and the Machine" (144) and in the Josiah Macy, Jr. Foundation conferences on what, from 1949 on, was referred to as "Cybernetics: Circular Causal and Feedback Mechanisms in Biological and Social Systems." It is beyond the scope of the present chapter to trace the future evolution (both good and bad) of work under the banner of cybernetics. Rather, let us simply note that as the field developed in the 50's, it began to fragment. Much work in cybernetics now deals with control problems in diverse fields of engineering, economics and the social sciences, while the broad field of computer science has become a discipline in its own right. Here we briefly cite four subdisciplines which have crystallized from the earlier concern with the integrated study of mind, brain and machine.

(i) Biological Control Theory: The techniques of control theory -- especially the use of linear approximations, feedback, and stability analysis -- came to be widely applied to the analysis of diverse physiological systems such as the stretch reflex, thermoregulation, and the control of the pupil.

(ii) Neural Modelling: The Hodgkin-Huxley analysis of the action potential, Rall's models of dendritic function, analysis of lateral inhibition in the retina, and the analysis of rhythm-generating networks are

examples of successful mathematical studies of single neurons, or of small or highly regular networks of neurons, which have developed in fruitful interaction with microelectrode studies.

(iii) Artificial Intelligence: This is a branch of computer science devoted to the study of techniques for constructing programs which will enable computers to exhibit aspects of intelligent behavior (101) -- such as playing checkers, solving logical puzzles or understanding (suitably restricted portions of) a natural language such as English. Buoyed by the slogan that "aeroplanes don't flap their wings", most practitioners of the field have turned their back on any concern with neural embodiments of intelligence. However, while some practitioners of AI (to use the standard abbreviation) look solely for contributions to technology, there are many who see their field as intimately related with cognitive psychology.

(iv) Cognitive Psychology: The concepts of cybernetics also gave rise to a new form of cognitive psychology which sought to explain human perception and problem-solving not in neural terms but rather in some intermediate level of information-processing constructs (100, 110). Recent years have seen strong interaction between AI and cognitive psychology.

Since the concern of cybernetics now extends far beyond the analysis of brain and machine, the term brain theory was introduced (6, 7, 8) to denote an approach to the study of brain rooted in the general perspectives of the 40's and 50's which sought to bridge the gap between studies of behavior and overall function (AI and cognitive psychology) and the study of physiologically and anatomically well-defined neural nets (biological control theory and neural modelling). The present chapter, then, is a perspective on past and potential contributions of brain theory

to the analysis of the role of visual information in the neural control of movement.

Concepts from Computer Science and Control Theory

To round out the introduction, we briefly review a few elementary concepts from computer science and control theory. The unifying notion is that of internal state -- whether it be the state of a control system or the state of execution of a computer program. In either case, we stress that, far from being a bundle of stimulus-response pairings, the behaving organism will act in a way which integrates its current sensory stimulation into a complex internal state which will be the true determinant of action. PROGRAMS NEED NOT BE STEREOTYPES. A program for an ordinary electronic computer executes one instruction at a time: to transfer inputs to memory, to combine pieces of data, to control output devices, and -- crucially -- to choose the next instruction on the basis of a test. Because of these tests the "overt behavior" of the program -- the temporal sequence of reading of inputs and emission of output values -- will depend on both the input values and data values already stored internally. This dependence yields the basic logical property of algorithms: a program can be specified in a form shorter than any normal execution upon particular data.

By contrast, normal English usage often takes "preprogrammed" as a synonym for "stereotyped", and for many neuroscientists, the word "program" is synonymous with the notion of a "fixed sequence". Consider, for example, the statements:

"We assume that the accuracy of reaching is an indication of the accuracy of the spatial calibration of the position cue for generating the program of reaching." (118, p. 275)

and

"A long-lasting sequence for a complicated action need not unroll in stereotyped fashion from its inception to the end: individual program segments can be fitted together intermittently, such as in tracking movements of the wrist allowing movements to evolve according to intermittent peripheral and central inputs." (23, p. 300)

For the computer scientist, it is the program that provides the formal description of the process which generates these different output sequences on the basis of differing values of inputs and internal parameters. To take a simple example, consider a man walking to the door. Depending on where he starts, his overt behavior might require 3 steps or 30 -- and each such sequence would constitute a "program" in the fixed-sequence sense. But, for us, it seems more insightful to hypothesize that a single program in the computer scientist's sense underlies all these behaviors. In one formalism, we might represent it as

do advance one step until door is reached.

[This can be recognized as a notational variant of the TOTE unit (Test-Operate-Test-Exit) (100).] Here, we explain all behaviors of the class in terms of a program with one action, advance one step, whose execution is repeated under the control of a single test, is door reached?.

If we regard such a program as a hypothesis about human behavior, we turn attention from the release of patterns with fixed numbers of steps to the study of ways in which perceptual mechanisms testing is door reached? may gate motor mechanisms. At a simple level, such "gating" differs little from negative feedback, but when we move to planning behavior in a complex environment, the loops within loops seem to call for the richer vocabulary that computer science can help provide. The development of such a vocabulary of "distributed control programs" will be a major goal of this chapter.

STATE AND FEEDBACK IN CONTROL THEORY. Basic control theory concepts were introduced in Chapter 10. Here we place those concepts in a slightly different perspective. In Chapter 10, the key concept was that of an operator S which converted the time history x of inputs to a system into the time history $y = S[x]$ of outputs of a system. However, we prefer an approach to system theory which makes the concept of state central. Consider a point mass undergoing rectilinear motion, so that at any time its position is given by $y(t)$ (the observable output of the system) and the force acting upon it is given by $x(t)$ (the input applicable to the system). To predict the response of this system we must not only know the mass m of the point (i.e. a point mass is a one-parameter family of models; a specific model is obtained when we fix the value of m) and the time history of forces applied to it, but also (according to Newton's laws) the initial position and velocity of the particle. We call the position-velocity pair at any time the instantaneous state of the system, and note that we can only associate a transfer function with a system if we fix the initial state (to equal the zero vector, for example).

More generally, then, a system can be expressed by a pair of equations

$$\dot{q}(t) = f(q(t), x(t)),$$

$$y(t) = \beta(q(t)).$$

The first expresses the rate of change $\dot{q}(t)$ of the state as a function of both the state $q(t)$ and the control vector $x(t)$ applied at any time t . The second equation expresses the current output $y(t)$ as a function of the current state $q(t)$. In the case of the point mass, β simply reads off the position vector from the combination of position and velocity vectors that constitute the state. In any case, we have that the transfer function of

any system will depend on the initial state, so that

$$y = S_{q(0)}(x)$$

will relate input to output for all times $t \geq 0$.

A control problem is to choose the input x in such a way as to cause the output y of a given system to behave in some desired way, whether to stay near a set reference value (the regulator problem), or to follow close upon some desired trajectory (the tracking problem). It is clear from the equation that the appropriate control signal to apply to the control system at any time can be computed on the basis of the instantaneous state of the system, if its exact value is available. But, in fact the controller can have at best an estimate of the state. Thus, the job of the controller can be divided into two parts, both of which must be accomplished in the face of unexpected disturbances.

(1) The estimation problem. To build up from the history of observations of the output of the controlled system a better and better estimate of the current state of the system.

(2) The control problem. To compute, on the basis of the estimate of the current state of the controlled system, a control signal that will guide the system appropriately.

A control signal defined by its intended effect may not achieve that effect either (as stressed in Chapter 10) because of the effect of disturbances upon the system, or because of inaccuracy in the controller's knowledge of the controlled system. Feedback is then required to compare actual and intended performance, so that a compensatory change in the input may be determined. Overcompensation yields instability; undercompensation yields poor adjustment to "noise". Thus, not only is feedback necessary, but it must be properly apportioned if the controller is to obtain smooth

coordinated behavior. In the next section we shall see how appropriate control systems may be thought of as "motor schemas", the basic units from which are built the coordinated control programs that control the organism's behavior.

SERIAL ORDER IN BEHAVIOR REVISITED. Many of Lashley's questions about "the problem of serial order in behavior" (90) are answered at the conceptual level (11) as soon as one thinks of the brain's computations not in terms of stimulus-response couplings or chains of associations but rather in terms of the above types of coordinated control programs. However, while our knowledge of computer programs removes the conceptual problem of serial order, the question of how such control strategies can be neurally implemented is only beginning to be answered. Much of the neurophysiological analysis of movement has focussed on spinal mechanisms, especially feedback mechanisms in posture and locomotion, and on higher level single-cell correlates of stimulus or response. Clearly, one of the aims of future research must be to better analyse the distribution of planning operations within cortical structures, and to understand the signal flow this planning must impose upon the cerebellum and other regions that modulate this planning.

VISUOMOTOR COORDINATION IN FROG AND TOAD

Many studies of eye movements represent the visual input by a few lumped variables representing, e.g., the angular position of a target. But if we wish to analyze the eye movements involved in reading or in scanning an outdoor scene, we must represent the usual input as a complex spatial structure. Visually guided locomotion, too, forces us away from the usual paradigms of biological control theory to stress that the input is a structured stimulus. In this section, we use an analysis of visuomotor coordination in frog and toad to suggest some of the ways in which structured visual stimuli may be processed by neural networks to control motor behavior.

Maps as Control Surfaces

First, some general background. A notable characteristic of brains is the orderly mapping from one neural layer to another, be it the retinotopic mapping from retina to the many visual systems, or the somatotopic mapping of motor cortex to the musculature. We briefly look at hypotheses as to how such a map may be viewed as a "control surface", so that it is spatiotemporal patterns in such a map that provide input to some control system in the brain.

A DISTRIBUTED PROCESSING MODEL OF SUPERIOR COLLICULUS. Julia Apter (3, 4) showed in the mid-40's that each half of the visual field of the cat maps topographically upon the contralateral colliculus. In addition to this

"sensory" map, she studied the "motor" map by strychninizing a single point on the collicular surface and flashing a diffuse light on the retina and observing which point in the visual field was affixed by the resultant change in gaze. She found that these "sensory" and "motor" maps were almost identical. This basic finding has, of course, been replicated and extended in many recent studies; but a review of such studies is beyond the scope of this chapter. What is important here is that Apter's studies were the basis for Pitts and McCulloch's (126) 1947 "distributed processing" model of the reflex arc from the eyes through the superior colliculus to the oculomotor nuclei to so control the muscles that direct the gaze as to bring the point of fixation to the center of gravity of distribution of brightness of the visual input. (With our current knowledge of retinal "preprocessing" we might now choose to substitute a term such as "general contour information", or any "feature" for "brightness" in the prescription above. But that does not affect the model that follows.)

Pitts and McCulloch noted that excitation at a point of the left colliculus corresponds to excitation from the right half of the visual field, and so should induce movement of the eye to the right. Gaze will be centered when excitation from the left is exactly balanced by excitation from the right. Their model (Figure 1) is then so arranged that, for example, each motoneuron controlling muscle fibers in the muscles which contract to move the eyeballs to the right should receive excitation summing the level of activity in a thin transverse strip of the left colliculus. This process provides all the excitation to muscles turning the eye to the right. Reciprocal inhibition by axonal collaterals from the nuclei of the antagonist eye muscles, which are excited similarly by the other colliculus, serve to perform subtraction. The computation of the quasi-center of gravity's

(Figure 1)




Figure 1. The Pitts-McCulloch scheme for reflex control of eye position via the superior colliculus: the eye can only be stationary when the activity in the two halves of the colliculus is balanced. (Adapted from W.H. Pitts and W.S. McCulloch, "How we know universals, the perception of auditory and visual forms", *Bulletin of Mathematical Biophysics* 9: 127-147, 1947. Reprinted with permission from Pergamon Press, Ltd.)

vertical coordinate is done similarly. Eye movement ceases when and only when the fixation point is the center of gravity.

COOPERATIVE COMPUTATION IN SOMATOTOPICALLY ORGANIZED NETWORKS. The above scheme shows how to design a retinotopically organized network in which there is no "executive neuron" that decrees which way the overall system behaves; rather, the dynamics of the effectors, with assistance from neuronal interactions, extracts the output trajectory from a population of neurons, none of which has more than local information as to which way the system should behave. In other words, the Pitts and McCulloch model of the superior colliculus showed how "the organism can be committed to an overall action by a population of neurons none of which had global information as to which action is appropriate."

Using the term "somatotopy" to include even retinotopic topographical arrangements, it has been argued (6) that the study of such cooperative computation in somatotopically organized networks provides a central paradigm in brain theory. For example, Boylls (20, pp. 181, 201-205) analyzes how the somatotpic relations between cerebellar cortex and related nuclei could provide a neural representation of the synergies involved in locomotion.

DISTRIBUTED MOTOR CONTROL. Here we wish to stress the way in which the position of activity in a layer of neurons may encode a control signal, so that the map provided by those neurons acts as a control surface for some control system. Where Pitts and McCulloch (126) modelled the output layers of superior colliculus as a control surface for eye movements, Braitenberg and Onesto (22) gave a scheme for the conversion of spatial input pattern

to a timing distribution of control signals for a ballistic movement. Such models led to the idea (5) that a plausible subsystem for vertebrate nervous systems may be of the type shown in Figure 2 in which position of the input on the control surface encodes the target to which the musculature will be sent. Further, we might expect that -- akin to the result of merging the Pitts-McCulloch scheme with the Braitenberg-Onesto scheme -- if an array of points is activated on the input surface, the system will move to the position which is the "center of gravity" of the positions encoded by that array.

The scheme of Figure 2 is just that -- specific further details are required to turn it into a testable model of a specific neural system. For example, in visual tracking, increasing angles of deviation might require movement of eyes alone, then of eyes and head, and then of eyes, head, and trunk. Thus the output of the motor-computer would not control a single joint but would control a whole hierarchy of subcontrollers, and would have to do so in a velocity-dependent way. Unfortunately, there is little successful modelling of specific controllers with the distributed structure of Figure 2. What we turn to, instead, are data and models on how structured visual stimuli might be processed to provide the appropriate input to such a controller.

(Figure 2)

A Model of Frog Snapping

Ingle (75) studied the snapping behavior of the frog when confronted with one or more fly-like stimuli. He found that in a certain region around the head, the presence of a fly-like stimulus will elicit a snap; that is, the frog will turn so that its midline is pointed at the "fly", and "zap" it with its tongue. When confronted with two "flies", either of




Figure 2. Schematic for a layered motor control system. A spatially-encoded target position is transformed into the appropriate sequence of motoneuron commands, with an array of inputs yielding movement to the "average" of the encoded targets. (From M.A. Arbib, The Metaphorical Brain, 1972. Reprinted with permission from John Wiley & Sons, Inc.)

which was vigorous enough that alone it would have elicited a snapping response, the frog could exhibit one of three reactions: it could snap at one of them, it could not snap at all, or it might snap at the "average fly".

A TECTAL MODEL OF CHOICE BEHAVIOR. We discuss Didday's model of such "choice behavior". This model is a simple one based on data available in 1970, and is presented not as state-of-the-art, but rather to provide a clear example of the processing of structured stimuli to provide the input to a motor controller of the kind posited in Figure 2. The task, then, was to design a network which can take a position-tagged array of "foodness" intensity from tectal neurons which modulate signals from the Group II "bug detector" cells of Lettvin et al. (93) with that from other layers and ensure that usually only one region of activity will influence the motor control systems. The model (33, 34) maintains the spatial distribution of information, with new circuitry introduced whereby different regions of the tectum so compete that in normal circumstances only the most active will provide an above-threshold input to the motor circuitry. To achieve this effect, we first introduce a new layer of cells in close correspondence to the "foodness layer", but whose activity is to yield the input to the motor circuitry. In some sense, then, it is to be "relative foodness" rather than "foodness" which describes the receptive field activity appropriate to a cell of this layer.

Didday's transformation scheme from "foodness" to "relative foodness" employs a population of what we shall call S-cells in topographic correspondence with the other layers. Each S-cell inhibits the activity that cells in its region of the "relative foodness layer" receive from the corresponding cells in the "foodness" layer by an amount that increases with increasing activity outside its region. This ensures that high

activity in a region of the foodness layer only "gets through" if the surrounding areas do not contain sufficiently high activity to block it.

When we examine the behavior of such a network, we find that plausible interconnection schemes yield the following properties:

- 1) If the activity in one region far exceeds the activity in any other region, then this region will eventually "overwhelm" all other regions, and the animal will snap at the space corresponding to it.
- 2) If two regions have sufficiently close activity, then:
 - (a) If both regions are very active they may both overwhelm the other regions and simultaneously "take command" with the result that the frog snaps between the regions.
 - (b) However, in many cases these two active regions will simply "turn down" each other's activity, and that in other regions, so much that neither is sufficient to "take command" and the frog will remain immobile, ignoring the two "flies".

One trouble with the circuitry as so far described is that the build-up of inhibition on the S-cells precludes the system's quick response to new stimuli. For example, in case 2(b) above, if one of those two very active regions were to suddenly become more active, then the deadlock should be broken quickly, but in the network so far described, the new activity cannot easily break through the inhibition built up on the S-cell in its region. In other words, there is hysteresis. Didday thus introduced what we shall call an N-cell for each S-cell. The job of an N-cell is to monitor temporal changes in the activity in its region. Should it detect a sufficiently dramatic increase in the region's activity, it then overrides the S-cell inhibition to enter the new level of activity into the relative foodness

layer. With this scheme, the inertia of the old model is overcome, and the system can respond rapidly to significant new stimuli. Didday hypothesized that the S-cells and N-cells modelled the sameness and newness cells, respectively, that had been observed in the frog tectum. Regrettably, no experiments have been done to test this hypothesis. We note here a number of other specific experiments suggested by the model:

(i) Record from a tectal T_5 unit (presumed to be the "decision unit" for orientation (67)) in paralyzed frogs, presenting a fly-like stimulus, A, centered in the ERF of the unit, and a second fly-like stimulus, B, in various positions. The model would predict an inhibitory zone for placement of B corresponding to the zone in which behavioral competition is observed; and would also predict a threshold effect, rather than simple additive inhibition, if the two areas are competing in a "winner-take-all" fashion.

(ii) Do the same experiment in frogs with a lesion of TP (thalamus/pretectum). If the inhibition is not purely intratectal (cf. the discussion of Ewert's experiments below) then the effect observed in the previous experiment should vanish or diminish.

(iii) Present the two fly-like stimuli in TP-ablated, behaving frogs. If the postulated lateral interaction is mediated by TP inhibition of tectum, then the frogs should always snap at the average fly, even for large separations.

(iv) Another experiment would test the hypothesis that prey-selection occurs at the tectal level rather than along the motor outflow. Present two unilateral fly-like stimuli to a behaving frog whose descending tectofugal pathways have been partially interrupted, with stimulus A mapping away from the interruption, while stimulus B corresponds to interrupted motor outflow. If competition takes place at the tectal-PT level, then

snapping at A should occur with the same frequency as in the normal, while there would be no response in those cases where a normal would respond to stimulus B. But if competition is downstream from the tectum, the animal should respond as a normal would respond to stimulus A presented alone.

Irrespective of the outcome of such experiments, the wealth of experimental data on visuomotor coordination calls for new models which extend the methodology of the Didday model, rather than incorporating the model directly. Such modelling is underway in my laboratory, and plans are being made for the interaction of theory and experiment. In the next subsection, we review some of the data the new models must encompass, but first place the Didday model within a broader context of neural modelling.

COMPETITION AND COOPERATION IN NEURAL NETS. The above model of prey-selection is an example of a broad class of models of what may be called competition and cooperation in neural nets. We provide a few pointers to papers on this topic. One of the first models of decision-making in neural circuitry to explicitly opt for cooperative computation was the S-RETIC model of Kilmer and McCulloch (86). In modelling the reticular formation of the brainstem, they used observations that the reticular formation helped switch the organism's gross state from sleep to wakefulness, and vice versa, to suggest that the function of the system was to determine the overall mode of behavior of the organism -- such as sleeping, fighting, or feeding. From the anatomical work of the Scheibels (133), they felt that a reasonable structural simplification of the system was as a stack of "poker chips". (This work of the Scheibels may be cited as one of the earliest contributions for "modules" of neural structure intermediate between brain region and neuron (105, 138, 139).) In the S-RETIC model, each module in the

stack receives a sample of the overall system input and, on the basis of that sampling, assigns weights to the different modes. However, the modules are coupled in such a way that each module readjusts its weights on the basis of activity in other modules to which it is connected. Kilmer and McCulloch were able to suggest a connection scheme which would lead to eventual consensus, with a majority of the modules assigning the greatest weight to a single mode. This is all done without any executive control.

We now turn to the problem of segmentation on depth cues. Julesz (84) designed "random-dot stereograms" in which each eye received a totally random pattern, but in which there were correlations between the inputs to the two eyes. Specifically, different regions in the two inputs were identical save for a shift in position, yielding a different disparity on the two retinas. Although such a pattern for a naive subject can initially appear to be nothing but visual noise, eventually disparity matching takes place and he perceives surfaces at different depths. Barlow, Blakemore, and Pettigrew (14) and Pettigrew, Nikara, and Bishop (123) found that cells in cat visual cortex are tuned for retinal disparity, and similar cells are posited in the human. Presumably, what causes the initial perception of noise is that, in addition to the correct correlation of points in the two retinas, there are many spurious correlations, and computation is required to reduce them.

The solution proposed by Dev (32) (see also (10, 97, 109)) was to imagine the cells of a given disparity as forming a population arrayed in a spatial map corresponding to the map of visual direction. Connections between cells were then to be arranged so that nearby cells of a given disparity would be mutually excitatory, whereas cells which were nearby in

visual direction but different in disparity would have inhibitory interaction. In this way, the activity of the array would become organized into a pattern in which, in each region of visual direction, cells of only one disparity type were highly active. As a result, the visual input would eventually be "segmented" into a number of distinct surfaces.

In the stereopsis model, then, we have competition along the disparity dimension and cooperation along the other dimensions. We note (7, 103) the similarity to S-RETIC, where the cooperation dimension is the row of modules, and competition is between modes rather than disparities. The Diddy model can be regarded as the limiting case where there is only a competition dimension, namely that of bug location. Such informal observations laid the basis for a rigorous mathematical analysis of competition and cooperation in neural nets (2). In terms of the possible use of Artificial Intelligence techniques in brain theory, we note that these neural mechanisms are similar to the "relaxation methods" for segmentation and region-labelling in machine vision (131, 143).

The Many Visual Systems

The prey-selection model provided a simple example of how a single task may be controlled in a distributed way, based on a structured stimulus array rather than a lumped input. Improvements on the model would incorporate data on pretectal involvement in snapping, would exploit the distributed structure of the model to analyze spatially-localized phenomena of both habituation and facilitation (40, 77), and would more fully explore the interaction of localization with the control of accommodation (25, 69, 75). In this section we shall review just a few of

these data which can enter into a neural model of how multiple controllers are coordinated in yielding an overall behavior.

COOPERATIVE COMPUTATION OF CONTROLLERS. In the simplest case of interaction of multiple controllers, the process of coordination simply "turns off" all but one of the controllers -- as when a frog either snaps at prey or flees a predator. In more complex situations, the activity of one controller will modulate the activity of another. Whether as experimentalists or modellers, our focus is on "chunks" of both brain and behavior which are sufficiently restricted to provide a coherent focus of investigation. Yet many of the properties of such "chunks" are governed by the rich interplay with other subsystems and involvement in other behaviors. Thus our models will have to be open-ended, able to interface with models of other sub-systems. Ingle (78) has argued for the utility of the frog in the analysis of such interactions. Five of the visual functions of the frog are prey-catching (a model for which we have already described), threat-avoidance, barrier negotiation, phototactic orientation and visual stabilization. In each case we may trace a different "visual map", providing a control surface for a distinct (yet not independent) layered motor controller. "Many types of visual representation of the world can be used to define the objects and spatial relations necessary for fine-tuning of a given motor sub-system." (78). Here is an example (76) of process coordination which should provide a fruitful basis for future experiments and modelling: Certain frogs preferentially respond to the approach of a large dark object (a "predator stimulus") by leaping away. For these "good avoiders", the jumping direction is highly predictable as a compromise between the forward direction and that directly away from the stimulus. Yet when a black stationary barrier is set within the frog's preferred jump path, the frog

escaping the "predator" will jump to one side or other of the barrier. The behavior is hardly surprising, but the challenge of unravelling the underlying neural mechanisms seems both tractable and important for the insight it can yield into the modification of stereotyped motor patterns by visual context.

PREY-ENEMY PATTERN RECOGNITION. Ewert has conducted a variety of behavioral, physiological and lesion experiments to determine aspects of the interaction between the tectum and other brain regions (especially pretectum and thalamus) in determining whether a moving object is treated by the toad as prey or enemy (see (40) for a review). Where normal toads will jump away from any large moving object, toads with thalamic-pretectal lesions (38) were uninhibited in their snapping at moving objects of all sizes, and would orient and snap at their own limbs, at another toad, or even at the experimenter's hand. A more analytic series of behavioral experiments (39) showed that elongation of a moving stimulus in the direction of movement increased the normal toad's prey-catching responses; whereas elongation orthogonal to the direction of movement decreased, and eventually inhibited, its prey-catching. However, in the case of the toad with thalamus and pretectum lesioned, orthogonal elongation yielded no such effect. These behavioral experiments have been complemented by physiological recordings (42 , 43) in which tectum type 2 neurons were seen to give responses to elongated rectangles quantitatively similar to the "prey-catching responses" measured behaviorally -- with the correspondence holding both in the normal toad and in the animal with pretectal/thalamic lesions. Moreover, small-field-units were found in PT (pretectum/thalamus) whose response increased with elongation of a rectangular stimulus orthogonal to the direction of motion.

These observations were the basis for a model of the prey-enemy recognition system (41) which used Fourier analysis to determine the transfer functions of the tectum type 2 cells and PT small-field-units so as to fit the experimental data (though not taking into account any size-constancy mechanisms that may exist). Their study looks at several alternative schemes for inhibitory interactions between the two "form filters" in generating orienting and avoiding responses.

The value of this model is its account of interaction between tectum and thalamus/pretectum in determining whether the toad will treat a single stimulus as "prey" or "enemy"; but the model can only handle a single stimulus. By contrast, the Didday model (admittedly for frog) addresses the problem of handling multiple stimuli, but treats the tectum in isolation. (The full thesis (33) does offer an ad hoc account of how features from different ganglion cells may be combined in determining "overall foodness".) The models are thus complementary, and should be subsumed in an analysis of the animal's behavior in complex structured environments.

Summary

We have shown in this section that the study of visuomotor coordination in frog and toad provides an area (by no means the only one) rich in experiments which may contribute to our understanding of perceptual structures and distributed motor control. We have briefly summarized a number of experiments on prey-catching and predator-avoidance, as well as on the modification of behavior by the presence of barriers. We have discussed models of prey-selection and prey-enemy pattern recognition. In addition, we have introduced the following ideas:

- 1) The need for techniques to analyze structured visual stimuli.
- 2) The concept of a neural map as a "control surface" which provides input to some control system in the brain.
- 3) The consequent notion of cooperative computation in somatotopically organized networks, with overall patterns of activity being generated in neural structures wherein no single neuron has global information as to what course of action is appropriate.
- 4) Within this paradigm, we noted the existence of a general theory of competition and cooperation in neural nets, with applications to prey-selection, mode selection and stereopsis, and with similarities to "relaxation techniques" used in Artificial Intelligence.
- 5) In addition to techniques for understanding competition and cooperation within the neural nets that constitute each controller, we need to understand the cooperative computation of controllers: most simply when coordination "turns off" all but one of the controllers; in general, when there is a rich interplay between subsystems each modulating the other.

PERCEPTUAL AND MOTOR SCHEMAS

The previous section emphasized simple problems of visuomotor coordination in frog and toad at a level of complexity appropriate for immediate interaction of theory and experiment. In this section and the next, we study concepts appropriate for the analysis of more complex behaviors. We shall see that a few of these concepts are well-related to experimental findings. Others appear logically necessary, but have not yet led to experiments in the neuroscience laboratory; and it may well require years of interaction between theorists and experimentalists to best adapt them to the analysis of neural processes.

Perceptual Schemas and the Action/Perception Cycle

THE INTERNAL MODEL. In most of the studies reported above on frog and toad, we could represent the environment in terms of a few stimuli to be classified as "prey" or "enemy". But when we turn to common human behavior, we know that our actions are determined by a far greater knowledge of the environment than afforded by the current stimulation of the retina -- especially since so little of that stimulation is foveal. Our actions are addressed not only to interacting with the environment in some instrumental way but also to exploring to update our "internal model of the world" (26, 61, 95, 101). In a new situation, we can recognize that familiar things are in new relationships, and use our knowledge of those individual things and our perception of those relationships to guide our behavior on that occasion. It thus seems reasonable to posit that the "internal model of the world" must be built of units each of which corresponds, roughly, to a domain of interaction -- which may be an object in the usual sense, an attention-riveting detail of an object, or some domain of social interaction. We

shall use the word (perceptual) schema to correspond to the unit of knowledge -- the internal representation of a domain of interaction -- within the brain (7).

(The notion of "schema" has been widely used in a variety of senses: in neurology (71, 48); in psychology (15, 108, 116, 125); (with such diverse names as 'frames' and 'scripts') in Artificial Intelligence (18, 102); and in the study of motor skills (134, 135).)

AFFORDANCES. The "ecological psychology" of Gibson (47, 52) downplays the hue-stimulus-intensity pattern of light falling on the retina, and instead stresses affordances (53) -- higher-level descriptors of substance and surface properties specific to the co-evolved relations of an animal in mutual interaction with its environment. We may identify these affordances with our perceptual schemas. Ecological psychologists stress the utility of describing the optic array in terms of those properties of the layout and dynamics of surfaces in the environment which, in relation to the environment of the animal, afford coordinated activity. (We shall devote a section below to a review of certain studies of "optic flow" based on pioneering work of J. J. Gibson.) Ecological psychologists claim that

"Meaningful information is detected directly (not constructed, not inferred, that is, not mediated epistemically) by an active and appropriately attuned animal. The sensitivity or attunement of the animal to affordances is wrought by the evolution of the species and by the individual animal's experience." (47, p. 7)

It may be epistemologically useful to insist that affordances rather than "raw sensations" are the building blocks of perception; but the neuroscientist must still analyze how neural circuitry can transform receptor activity (itself modulated by efferent control of receptors, and selected by exploratory and other motions of the organism) into a continually updated

structure of affordances.

(Figure 3 about here)

THE ACTION-PERCEPTION CYCLE. To better appreciate the intimate relation between perception and action, consider Figure 3. The inner cycle is referred to as the perceptual cycle (108). The subject's exploration of the visual world is directed by anticipatory schemas, which Neisser defines to be plans for perceptual action as well as readiness for particular kinds of optical structure. The information thus picked up modifies the perceiver's anticipations of certain kinds of information which, thus modified, direct further exploration and become ready for more information. For example, to tell whether or not any coffee is left in a cup we may reach out and tilt the cup to make the interior visible -- and keep tilting the cup further and further as we fail to see any coffee until we either see the coffee at last or conclude that the cup is empty. (For an expectation mechanism modelled at the level of neural nets, see (64); for a neural net model of reorganization of a developing sensory or motor code based on expectation mismatch, see (66).) We here stress that one cannot understand perception unless it is embedded within the organism's on-going interaction with its environment (see also (6)). For (consider the outer cycle of Figure 3) as the organism moves in a complex environment -- making, executing and updating plans as it does so -- it must stay tuned to its spatial relationship with its immediate environment, anticipating facets of the environment before they come into view. The information gathered during ego-motion must be systematically integrated into its internal model of the world (which may also be referred to as a cognitive map (140), which is not so much a mental picture of the environment as an active, information-seeking process.




Figure 3. The Action-Perception Cycle. (From U. Neisser, Cognition and Reality: Principles and Implications of Cognitive Psychology, 1976. Reprinted with permission from W.H. Freeman and Company and the author.)

PLANNING. The intelligent organism does not so much respond to stimuli as it selects information which will help it achieve current goals -- though a well-designed or evolved system will certainly need to take appropriate account of unexpected changes in its environment. Planning is the process whereby the system combines an array of relevant knowledge to determine a course of action suited to current goals. In its fullest subtlety, planning can involve the refinement of knowledge structures and goal structures, as well as action *per se*.

Novel inputs (e.g. coming upon an unexpected obstacle) can alter the elaboration of high-level structures into lower-level tests and actions which in turn call upon the interaction of motor and sensory systems. We seek to study programs which are part of the internal state of the system, and which can flexibly guide ongoing action in terms of internal goals or drives and external circumstances.

PERCEPTION AS POTENTIAL ACTION. Our thesis is that perception of an object (activating appropriate perceptual schemas) involves gaining access to routines for interaction with it (motor schemas), but does not necessarily involve execution of even one of these routines. While an animal may perceive many aspects of its environment, only a few of these can at any time become the primary locus of interaction. We may say that perception activates (defining a search space; drawing a map), while planning concentrates (laying out the route to be followed).

COORDINATIVE STRUCTURES. We know (62, 136) that we may view the basic pattern of regular locomotion as generated by spinal oscillators tuned by inter-limb coordination, subject in turn to modulation by brainstem and other structures -- a "coordinative structure" in which the control of movement does not proceed through the independent control of each motor

unit. (For a use of these concepts in a model of cerebellum, see (20).) Noting this, Fitch and Turvey (47, p. 16) argue against the view that the organism needs plans of behavior which "might be said to represent symbolically and discretely each of the segments [of a behavior] in their appropriate order." Rather, they note that,

"[while] it is true that quadruped locomotion can be (a) partitioned by eye into segments that repeat in a different order and (b) described kinematically in terms of distances, directions, velocities, accelerations, etc., there is no reason to believe that either the segments and their order or the kinematic details are symbolically represented -- anywhere."

Thus, in contrast to the notion of plan, they argue (47, pp. 24-5) for "the arising of constraints" rather than "the issuing of commands".

"We may envision an evolving ecosystem in which those subsystems (structure, constraints) arise which share the greatest mutual compatibilities with the whole system. ... Change does not happen by an agent monitoring the progress of an evolving subsystem and issuing corrective commands to hurry it along toward an intended goal; change occurs by the mutual fitting together of the simultaneously changing subsystems."

The ecological psychologist is right to stress that "change occurs by the mutual fitting together of the simultaneously changing subsystems". What we must ask, though, is what are the actual mechanisms for coordination of a variety of control structures. The neuroscientist must analyze the neural embodiments of such coordinations, and explain not only the impairment due to brain damage but also the ability to mobilize different resources to behave in the face of such disturbances (74). We hypothesize that a level of "planning" to generate "plans" or "coordinated control programs" is a necessary overlay for the control system/coordinative structure analysis in any fully developed theory of motor control, and suspect that this is in reality consistent with the "ecological" approach once the need to spell out detailed mechanisms is admitted.

Neural data make no sense except in context. The "coordinative structures" and the action-perception cycle remind the neuroscientist to eschew too Sherringtonian a view of a brain responding to an environment via a chain of reflexes (90).

THE SCHEMA-ASSEMBLAGE. We have used the term "schema" for the process whereby the system determines whether a given "domain of interaction" is present in the environment. The state of activation of the schema will then determine the credibility of the hypothesis that what the schema represents is indeed present; while other schema parameters will represent properties such as size, location and motion of the perceived object.

Consider a schema that represents, say, a chair; and consider an environment that has two chairs in plain view. It is clear that two copies of the chair-schema -- or, at least, two separate sets of chair-schema-parameters -- will be required to represent the two chairs. We refer to these two copies as being separate "instantiations" of the same schema, each with its own set of parameter values. We may thus view the internal representation of the environment as an assemblage of spatially-tagged, parametrized, schema instantiations.

EYE-MOVEMENTS AND VISUAL PERCEPTION. To indicate how a bridge may be built between neural net studies and the present rather abstract considerations, we briefly recall a study which linked frog prey-selection to human control of eye-movements during visual perception. The study is based on the hypothesis that since frog tectum is homologous to mammalian superior colliculus, the transformations within frog tectum may yield insight into the transformations taking place in the superior colliculus of monkeys, and of humans. Where suitable visual input in the frog triggers snapping at an object, the visual input to the superior colliculus can trigger an

eye movement, as when we look at the source of a sudden movement or a flash of light.

(Figure 4 about here)

Noting this, Didday and Arbib (35) modelled the role of eye movements in visual perception. The neocortex is modelled as a system which constructs and updates the schema-assemblage. Lumped modelling omits most of the crucial structure in analyzing what is "fovea-worthy" in a structured visual pattern. The superior colliculus is modelled as Didday (34) modelled the frog tectum. Thus the neocortex is not modelled at neural network level, although the superior colliculus is. In the model (Figure 4), the superior colliculus not only receives the direct input concerning the sudden flashes of light, etc., which can distract attention, but also receives "calls" from the neocortex, as when a schema needs more detailed input from the fovea to confirm a hypothesis about the presence of an object it represents in a particular region; or when there seems to be a discrepancy between what the schema has predicted and what is actually happening in a region of periphery, which thus demands more scrupulous attention (cf. the earlier discussion of anticipatory schemas).

This model can account for the scanpaths -- systematic paths followed by the eyes in repeated viewing of a picture -- seen by Noton and Stark (112). These authors explain scanpaths by positing that the brain stores the picture as a list of alternative visual features and the eye movements required to move from one feature to the next. As an alternative, Didday and Arbib suggest that the brain combines current input with the pattern of schema activation to determine each eye movement by a highly parallel computation, rather than by relying on memory of previous movements.

The idea of a hybrid model such as that of Figure 4 (in which only certain subsystems have been refined to the neural level) is that one can

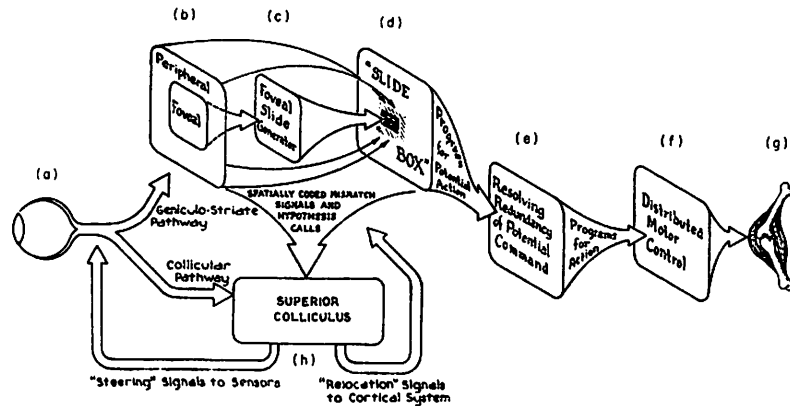


Figure 4. The general form of the model, showing the relationships between the functions of the midbrain and the cortical visual systems, the "slide-box" and the motor pathways. The "slides" of their model correspond to the "schemas" of the present discussion, with the state of the "slide box" corresponding to the notion of a "schema-assemblage". (From R.L. Didday and M.A. Arbib, "Eye movements and visual perception: a 'two visual system' model", *Int. J. Man-Mach. Stud.* 7: 547-569, 1975. Reprinted with permission from Academic Press Inc. (London) Ltd. and the authors.)

hope to use the posited interaction to suggest experiments upon the signal flow between brain regions; and that, to the extent that these experiments succeed, one can then more tightly constrain the model to bring each subsystem closer to a neural net analysis.

Optic Flow and the Control of Movement

A number of studies of the Gibsonian notion of affordance have focussed on optic flow -- the pattern of motion of optic elements on the retina induced by relative motion of organism and environment -- and the information it can provide for the behavior of the organism (50, 51, 54, 57, 91, 92, 127). In this section, we give a brief introduction to how the optic flow can provide exproprioceptive information (92) about the position, orientation and overall movement of the organism relative to its environment. The development is a variant of that of Lee (91) in that we use a planar retina rather than a cylindrical one. We shall analyze only monocular information, ignoring depth cues afforded by stereopsis.

FOCUS OF EXPANSION. Let ξ be the horizontal coordinate and η the vertical coordinate of a point on our planar retina. The organism-centered spatial coordinates are then x and y corresponding to ξ and η , respectively, with z the horizontal coordinate in the organism's line of gaze (Figure 5).

(Figure 5)

The gaze of the organism is not restricted to coincide with the direction of locomotion. We may express the velocity of the organism as a vector (u, v, w) . In the current study, we shall assume $v = 0$ (no vertical motion -- smooth passage over a perfectly flat terrain) but do not assume $u = 0$ (i.e. we do not restrict motion to be along the direction of gaze). Note that the retinal display -- and thus the optic flow derived from it --




Figure 5. To simplify visualization of the correspondence between retinal and spatial coordinates, we have placed the focal point behind the retina -- but this is, of course, equivalent to placing the retina behind the focal point, and then inverting the coordinates. We thus have $\xi = ax/z$, $\eta = by/z$ where a and b are positive scale constants.

is based only on relative coordinates. Thus the same flow will result from the organism moving with some velocity in a stationary environment; or from the environment as a whole moving with the opposite velocity about a stationary organism.

(Figure 6)

Just as when driving along a straight road we see all points of the landscape appearing to radiate from the horizon point on the road, so it can be shown that for the relative motion of Figure 6, all the trajectories of the corresponding optic flow (Figure 7) will emanate from a single point

(Figure 7 about here)

$(au/w, 0)$ called the focus of expansion (FOE). The FOE is the retinal projection of the horizon point toward which the organism is locomoting.

Each flow line in Figure 7 is the projection on the retina of some texture element in the environment, as plotted over some period of time. Each point of an object will describe its own flow line, and it is clear that the orientation of the flow line has no relation to the shape of the object.

RELATIVE MOTION. The pattern of optic flow depends only on the velocity of texture elements relative to the organism, whether it is the organism, environment or both that is moving. (The organism will, of course, have other information available about its own motion, such as knowledge of its motor activity, to help it compute what a particular pattern of relative motion means.) If the organism is progressing through an environment, and a single object is moving relative to that environment, the points on that object have a velocity relative to the organism different from that of the background. We will thus see two different FOEs, one for the flow generated by the object, and one for the flow generated by the background (Figure 8).

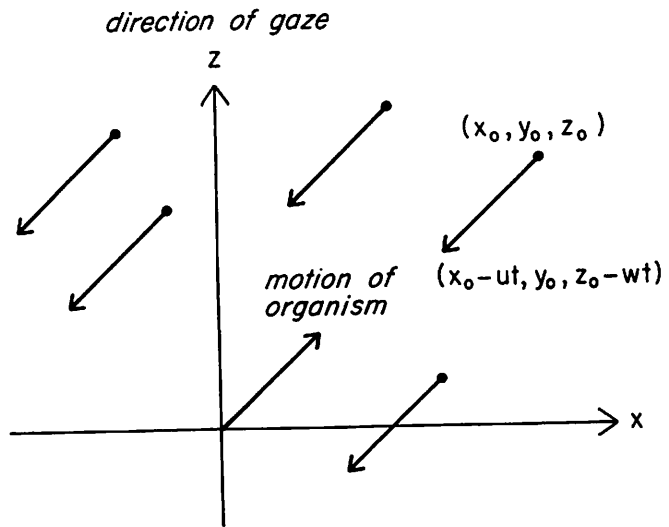


Figure 6. Relative motion of organism and environment.




Figure 7. Optic flow radiates from a common Focus of Expansion (FOE) when the motion of organism relative to the environment is constant and forward.

(Figure 8 about here)

INFERRING COLLISION INFORMATION. The crucial point about the optic flow, in terms of its role as a Gibsonian affordance, is that it contains information which may be useful to the organism in guiding its locomotion. For example, the optic flow contains enough information to tell the organism whether or not it is on a collision course with a texture element moving with constant relative velocity. The display in Figure 8 is useful, because the variable distance between arrowheads reminds us that each optical element has an associated retinal velocity as part of its characterization in the optic flow. If the arrow is pointing up and to the right, say, then the corresponding texture element is moving up and to the right relative to the center of the retina. The question is -- since the organism has finite extent, will the texture element move far enough before drawing level with the organism to avoid collision.

Suppose the organism has width $2k$. Then, in the case in which the direction of locomotion coincides with the direction of gaze so that x is constant, there will be eventual collision of the organism with the texture element at (x_0, y_0, z_0) just in case $|x_0| < k$. Now we know from Figure 7 that when $u = 0$ and $w > 0$, the flow lines radiate out from the FOE with coordinates $(0,0)$. As we may see from Figure 9, the closer the texture

(Figure 9 about here)

element, the larger is its $\dot{\xi}(t)$ for a given $\xi(t)$. That is, for a given visual angle, the further the texture element, the smaller its associated flow. The smaller $\dot{\xi}(t)$, then, the less likely is a collision. In fact, it can be shown that collision occurs if $|\dot{\xi}(t)| > \left| \frac{w\xi(t)^2}{ak} \right|$. For completeness, we note that collision will not occur if the object is one like an overhanging tree which can pass over the organism's "head". In fact,




Figure 8. Computer output showing two superimposed optic flows each radiating from its own FOE. One is due to the forward progression of the organism; the other due to the motion of an object moving toward the organism. The tail of each arrow shows the initial projection of a texture element on the retina; the three arrowheads indicate the retinal projection of the same texture element after each of three subsequent steps.




Figure 9. For a given $\xi(t)$, the closer the texture element the larger is $|\dot{\xi}(t)|$.

collision will be avoided if $0 < \dot{\eta}(t) < \frac{wn(t)^2}{bh}$, where h is the height of the top of the organism above its "eye".

Note that the above expressions require use of the ratios w/ak and w/bh . We might hypothesize that ak and bh are stored within the organism (perhaps as a result of adaptive learning) while w will vary from occasion to occasion and so must be inferred from motor commands, kinesthetic information or by visual inference using, e.g., targets of known size for calibration.

TIME UNTIL ADJACENCY. It is also possible to infer from the optic flow how long it will be until a texture element will draw adjacent to the organism (i.e. how long until its z -coordinate drops to zero). In the particular case in which the element is on a collision course, the time until adjacency is in fact the time until contact should the relative motion of organism and object remain unchanged. It thus affords the organism the information it needs to initiate coordinated action should contact be planned (as when a bird extends its legs when landing); or to change course at an appropriate speed if it plans to avoid contact.

Motor Schemas

We have used the term "schema" to indicate the type of unit from which the internal representation of the environment may be built; and we have used optic flow to indicate the type of computation whereby action-relevant parameters may be extracted from the dynamic sensory input. In this section, we shall suggest that programs for motor control may themselves be seen as assembled from suitable units, which we shall refer to as "motor schemas". We first relate these motor schemas to synergies in the sense of the Russian school founded by Bernstein; and then, after a

brief review of relevant control theory, we look at experimental evidence for the adaptability of these motor schemas.

SYNERGIES. The first concept of synergy, due to Sherrington, posits a reflex unit above that of the motor unit; while the second, due to Bernstein, suggests that a restricted number of programs may underlie most of our behavior. The Sherringtonian concept of synergy may be seen in the following quotations:

"The executant musculature ... provides a reflex means of supporting or reinforcing the co-operation of flexors with flexors, extensors with extensors, etc. The proprioceptors of reaching muscles operate reflexly upon other muscles of near functional relation to themselves. Active contraction (including active stretch) and passive stretch in the reaching muscles are stimuli for reflexes influencing other muscles, and the reflex influence so exerted is on some muscles excitatory and on others inhibitory; it is largely reciprocally distributed, knitting synergists together." (31, p. 129)

"A muscle which by fixing a joint enhances the effect of another muscle crossing that joint to act on a further one, is said to be the latter muscle's pseud-antagonist. ... Pseud-antagonism is really a form of synergism and reflex co-ordination deals with it as with other synergism not by reciprocal innervation but by co-contraction." (31, p. 131)

Thus, for Sherrington, the synergy is an anatomically based reflex linkage of a group of muscles. By contrast, Bernstein's concept of synergy, which is the key to Russian studies of motor control, is a functional one, a dynamic pattern changing over time during some motor act:

"[A] complex synergy is involved in walking. ... 'The biodynamic tissue' of live movements [appears] to be full of an enormous number of regular and stable details." (17, p. 67)

"[In old people] the synergy existing in normal walking between the action of the arms and legs is destroyed." (17, p. 93)

Later work of the Moscow school came to view synergies as the functional building blocks from which most motions can be composed:

"In order for the higher levels of the central nervous system to effectively solve the task of organizing motor acts within a required time, it is necessary that the number of controlled parameters be not too large, and the afferentation, requiring analysis, not too great.

[This is achieved] by the so-called synergies. ... Each synergy is associated with certain peculiar connections imposed on some muscle groups, a subdivision of all the participant muscles into a small number of related groups. Due to this fact, to perform motion it is sufficient to control a small number of independent parameters, even though the number of muscles participating in the movement may be large. Although synergies are few in number, they make it possible to encompass almost all the diversity of arbitrary motions. One can separate relatively simple synergies of pose control (synergy of stabilization), cyclic locomotive synergies (walking, running, swimming, etc.), synergies of throwing, striking, jumping, and a certain (small) number of others." (49, p. 162)

We shall restrict the usage of "synergy" to the Sherringtonian sense, while using the name "motor schema" for the Bernsteinian sense. Having said this, we may note that the motor schemas described in the last quotation -- such as walking and running -- are at too high a level for a direct linkage to neural analysis, and that hierarchical refinement is required: be it of walking into coordinated stepping patterns, or of writing into letter formation and thence into the formation of specific strokes under visual and tactile control. These schemas, at whatever level of refinement, then correspond to the individual control systems in the coordinated control programs that we shall describe below. (For an earlier theoretical perspective on these matters, see (58, 59, 60).)

SYSTEM ADAPTATION. Having suggested that motor schemas be viewed as control systems, we now review certain concepts from control theory which help us analyze how such schemas may exhibit "short-term memory". To properly control the motion of an object (the "controlled system"), the controller must clearly know relevant parameters of the object, such as its mass and movement of inertia. However, in the real world the exact values of these parameters are seldom available, and may actually change over time (compare short-term loading effects on muscles and longer-term aging effects and weight changes). To adapt to such changes the outer, feedback, loop of

Figure 10 must be augmented by an "identification algorithm". The job of this algorithm is to more accurately identify the parameters of the controlled system. To do this, it continually monitors the behavior of the controlled system and compares it with the output that would be expected on the basis of the current estimated parameters. Any discrepancies in the output can be used to obtain more accurate estimates of the parameters that

(Figure 10 about here)

define the controlled system. These updated parameters can then be supplied to the controller as the basis for its state estimation and control computations. (For an example of an identification algorithm specifically developed for a controller of limb movements, see (128). Parametric equations of motion were updated via a quantized multi-dimensional memory on the basis of analysis of practice movements. Tests with a small computer and mechanical arm revealed an ability to acquire new movements, and adapt to changing loads.)

If the controlled system or the disturbances to it are sufficiently slowly time-varying for the identification procedure to make accurate estimates of (system + disturbance) parameters more quickly than they actually change, the controller will be able to act efficiently, despite the fluctuations in controlled system dynamics. The controller coupled to an identification procedure is an "adaptive controller": it adapts its control strategy to changing estimates of the dynamics of the controlled system. Houk in Chapter 10 presents a careful analysis of that class of adaptive controllers that operate explicitly by "model reference". To this we simply add that Ito's (78, 79, 80, 81) ideas on feedforward in vestibular-cerebellar interactions are based on this concept, as well as on the synapse-adjustment cerebellar model of Marr (96). (For other approaches to cerebellar plasticity see (55, 56, 63, 120, 121, 122). For the related topic

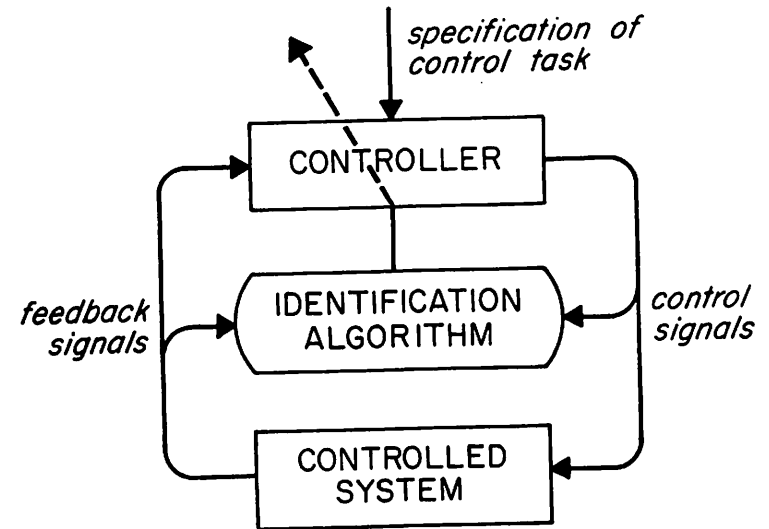


Figure 10. To render a controller adaptive, an identification algorithm monitors control signals and feedback signals to provide the controller with updated estimates of the parameters that describe the controlled system.

of learned motor expectancies see (64, 65, 66).)

Minsky (101) has observed that it may also be necessary to have the identification procedure generate some of the input to the controlled system -- in other words, apply test signals to try out various hypotheses about the parameters of the controlled system -- trading off the degradation of control due to an inaccurate estimate of the state parameters against the degradation due to having the controller intermittently relinquish control to the identification procedure. Allum (1) studied human response to displacements acting to rotate the shoulder while the subject was trying to maintain constant position against a pre-existing force. Of the four stages of the response, it is the second stage that is often considered to be a stretch reflex. But Allum noted that it could only provide about 15% of the force required, and so hypothesized that this was in fact a test signal, and that the resultant information about the current loading of the muscle was used by the CNS to generate the third stage of the response, a "long latency" increase in force which restored the arm to its original position.

The above discussion emphasizes the "short-term memory" of the schema, updating parameters relevant to the current controlled system. But an identification algorithm may also have a "long-term memory" which determines parameters relevant to the controller's general control strategy.

Note well that the identification algorithm can only do its job if the controller is of the right general class. It is unlikely that a controller adapted for guiding the arm during ball-catching will be able, simply as a result of parameter-adjustment, to properly control the legs in the performance of a waltz. Thus the adaptive controller of Figure 10 (controller + identification procedure) is not to be thought of as a model of the brain; rather each such control system is a model of a single motor schema which

can be activated when appropriate. The language of "coordinated control programs" presented below addresses the problem of the coordinated phasing in and out of the brain's manifold control systems.

MUTABILITY OF MOTOR SCHEMAS. Motor schemas simplify motor control by restricting the degrees of freedom of the system. Were such generators fixed, many possible movements would be eliminated which are *a priori* well within the organism's muscular repertoire. However, in man at least, the schemas can change, adapting with a time course far slower than that of the movements they control.

A study of microstimulation of inferior olive in locomoting premammillary cats (21) showed that many odd as well as straightforward combinations of muscles were depressed or potentiated. More importantly, alteration of muscle activity lasted for many seconds after stimulation, during which time the cat would take quite a few steps using these new muscle relationships. This seems to be a short-term memory which may be consolidated only by continued use of (or failure to "undo") the altered pattern of muscle coupling. Cerebellar involvement in such short-term adjustments was predicted by Boylls' (19, 20) model of the cerebellum. This study marshalled a wealth of anatomical and physiological data to model how the spino-olivo-cerebellar climbing fiber apparatus of the cat cerebellar anterior lobe could, in principle, exert prolonged "tonic" biases upon muscle usage during locomotion. An interesting feature of the model is a predicted association between the Oscarsson climbing fiber strips in cerebellar cortex and various "structural" muscle biases associated with different gaits. The model showed how the compartmental organization of the anterior lobe and its matching Oscarsson strips (117) might enable the anterior lobe to telescope the changed relationships with a temporal sequence of movements into a

compact spatial representation, based on somatotopic relationships between the anterior lobe and its input and output nuclei.

Another example of mutability is afforded by Pal'tsev's (119) demonstration of human subjects' unconscious acquisition of a new "spinal reflex organization". The subject sits on the edge of a table with legs hanging free. A conditioning patellar tap given to one knee of a naive subject would reduce the patellar reflex in the leg for a short period of time. Pal'tsev attributed this to reflex circuitry similar to that found in spinal cats. However, if a series of about 20 tendon taps were given simultaneously to both knees, a subsequent conditioning tap to one would augment the test reflex in the opposite leg. This effect also followed the subject's simultaneous extension of both knees a number of times. In either case, the newly acquired reflex structure would persist for several minutes.

This process of adaptation of motor schemas need not be under conscious control. For example, Nashner (107) found that subjects standing on a rotating platform may make inappropriate "reflex" responses, in terms of stabilizing posture, the first few times the platform rotates. Subsequently, the "reflex" adapts.

Summary

A framework is offered for analyzing visually-guided behavior of a complex organism. The framework is based on the following general premisses:

1) The action/perception cycle: As the organism moves -- making, executing and updating plans -- it must maintain an up-to-date representation of its spatial relationship with its environment.

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

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

4) The plan of action is to be thought of as a program composed of units called "motor schemas" (also referred to in the Russian literature as "synergies"). Each motor schema may be viewed as an adaptive controller which incorporates an "identification procedure" to update its representation of that which is being controlled. Thus the identification procedure may be viewed as a perceptual schema embedded within a motor schema. The next section will discuss the possible nature of "coordinated control programs" which can coordinate the activation of motor schemas.

The presentation of this framework was augmented by an exposition of optic flow and the information it can provide for controlling the organism's behavior, as an example of the Gibsonian notion of affordance/schema; and a review of experiments indicating the adaptability of motor schemas.

COORDINATED CONTROL PROGRAMS

Biological control theory usually studies neural circuitry specialized for the control of a specific function, be it the stretch reflex or the vestibulo-ocular reflex. Yet most behavior involves complex sequences of coordinated activity of a number of such control systems. We will thus explore the notion of a "coordinated control program" as a combination of control theory and the computer scientist's notion of a program suited to the analysis of the control of movement. Control theorists use a block diagram to represent a system: each box represents a subsystem which is continually active; while the lines linking the boxes transfer data, showing how the output of one system helps determine the input to another. By contrast, the boxes of the flow diagrams used by computer scientists to represent not subsystems but patterns of activation of subsystems. The computer has various subsystems such as memory registers, arithmetic units and test units. At any time in a computation, certain data will be stored in these subsystems, and one box of the flow diagram will be "activated" in the sense that it will be used by the computer to determine what tests and operations will be carried out by the subsystems and how data are to be transferred amongst them. The lines of the flow diagram then specify how activation is to be transferred from one instruction to another.

Below, we analyze motor performance in terms of coordinated control programs which can control the time-varying interaction of a variety of control systems. In the diagrams representing such a program, there will be lines representing both transfer of activation and transfer of data. But first we devote a subsection to a brief recapitulation of concepts from biological control theory.

Feedforward

In Chapter 10, Houk has defined feedforward (58, 95) as that strategy whereby a controller monitors disturbances to a system directly and immediately applies appropriate compensatory signals to the controlled system, rather than waiting for feedback on how the disturbances have affected the system. The advantage is speed, disturbances may be compensated before they have any noticeable effect on the system, but the cost is paid in controller complexity: to "pre-empt" the disturbances, the controller must have an accurate model of the effect of disturbances upon the system, and not just of the system's response to "normal" inputs. Here we extend the sense of feedforward to include a strategy that generates large control signals to rapidly bridge large discrepancies in desired output at too great a velocity for long-latency feedback paths to play a major effect.

DISCRETE-ACTIVATION FEEDFORWARD. In the hypothetical scheme of Figure 11 (which does not show the identification algorithms which may provide adaptive components), feedforward is activated only when the error is not small. If well-calibrated, the feedforward controller will, with a single brief time-pattern of control return the system to the "right ballpark", i.e. making the error small enough for feedback control to function effectively. The system should thus have a "refractory period" based on the time-constants of the controlled system -- it should not generate a second control signal before the control system has had time to respond fully to the first control signal.

The sample of the system's output is called "feedback" when fed to the feedback controller in Figure 11, yet is called "actual output sample" when fed to the feedforward controller. This looks like a semantic trick, but

(Figure 11)

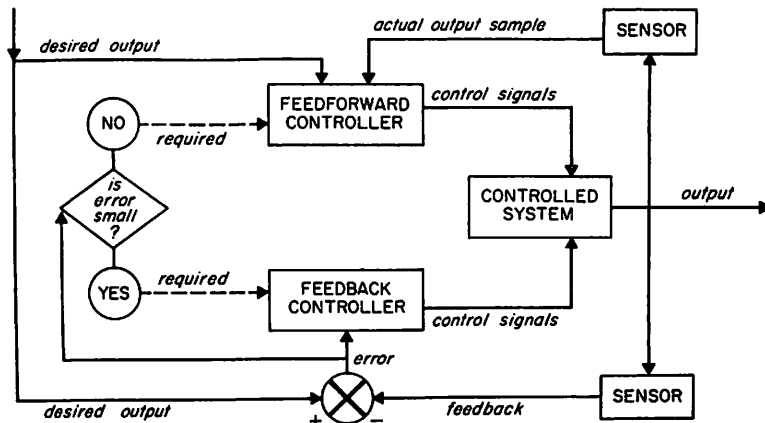


Figure 11. Discrete-activation feedforward -- one of various possible configurations in which feedback and feedforward controls are explicitly separated. Here feedforward is active for large errors to get the controlled system "into the right ballpark", while feedback provides "fine-tuning" in the presence of small errors. The dashed lines marked "required" indicate the supply of necessary activation if the system supplied is to function. Non-dashed lines indicate "data flow".

the difference is, in fact, genuine. A feedforward controller will, in general, need to know the actual state of the controlled system before each discrete activation but need not monitor that output while the control signal is actually emitted. By contrast, the feedback controller continually monitors the error signal in generating its controls.

In Figure 11, we have shown the feedback controller as active only when the error is small. But if, in fact, its action is slow and small compared to that of the feedforward controller (also called an "open-loop controller" in this context) there may be no need to disconnect the feedback controller. There does not seem to be any experimental analysis of the switching posited in the diagram.

The situation in Figure 11 might be refined to have the feedforward controller monitor the relation between the actual trajectory and a predicted trajectory, changing strategy if the discrepancy exceeds a threshold. But, again, we have a discrete-activation form of feedforward.

CO-ACTIVATION FEEDFORWARD. The spring model of muscle (44, 73) shows that we must consider control strategies in which a functional relation exists between a desired output (e.g. muscle length) and a necessary input (e.g. maintained motoneuron firing). In this case, the feedforward would be co-activated with the feedback system as shown in Figure 12 providing the approximate control signal required to maintain the desired output, with feedback serving to refine this approximation.

We may view eye-movement control as an example of the discrete-activation mode (a saccade being effected by discrete-activation feedforward; while feedback controls tracking); while we may view α - γ co-activation as an example of the co-activation mode (with the α -system feedforward being co-activated with the γ -system feedback).

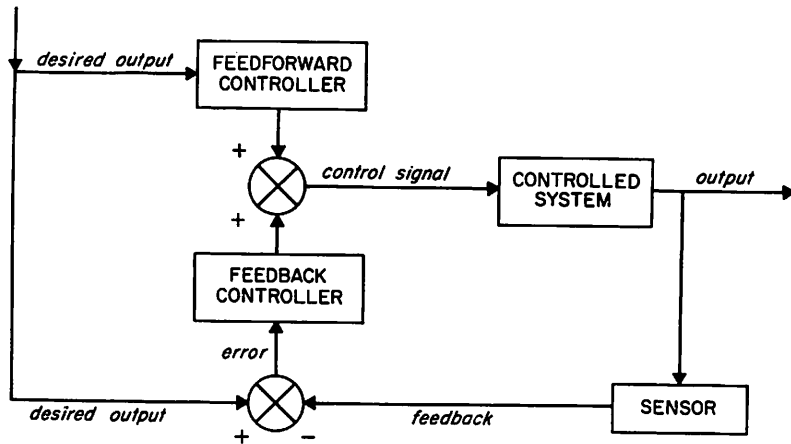


Figure 12. Co-activation feedforward -- one of various possible configurations in which feedback and feedforward are explicitly separated. Here the feedforward controller continually supplies a control signal which can maintain the output of the controlled system "in the right ballpark", while the feedback controller utilizes error feedback to provide the necessary fine-tuning to compensate for inaccuracy in the feedforward controller's model of the controlled system, as well as for disturbances. Such a mode of control is appropriate only when the controlled system has a functional relation between maintained input and maintained output:

MOTOR SCHEMAS AGAIN. The concept of schema has developed a special meaning in the motor skills literature, e.g. in the work of Schmidt (134, 135). Each of Schmidt's schemas (cf. Figure 13) controls the performance of a single motion in the laboratory or in sports (e.g. swinging a bat) rather than complex manipulation or goal-oriented performance in a dynamic environment. The schema is broken into two parts: The recall schema seems akin to the feedforward of Figure 11, being responsible for the complete control of a rapid movement, even though environmental feedback may later signal errors. The recognition schema is responsible for the evaluation of response-produced feedback. It seems to combine the on-line feedback of Figure 12 and the identification procedures of Figure 10 which may operate even after a movement is completed to better tune the schema for its next activation.

(Figure 13)

Interwoven Activation of Motor Schemas

A "sequence" of actions may be quite complex, with the actions interwoven and overlapping. Simultaneous actions must be coordinated, and successive actions must be smoothly phased one into the next. In this subsection, we discuss the concept of a "coordinated control program" as the type of structure which could orchestrate the interwoven activation of the motor schemas which control the different actions.

A PROGRAM FOR GRASPING. Consider, first, a human reaching to grasp an object. We have a ballistic movement towards the target during which the fingers are adjusted to the size of the object, and the hand is rotated to the correct orientation. When the hand is near the object, a final feedback adjustment is made in the position of the hand. A hypothetical "coordinated control program" for this behavior is shown in Figure 14: the solid arrows indicate transfer of data; broken arrows convey "activation signals".

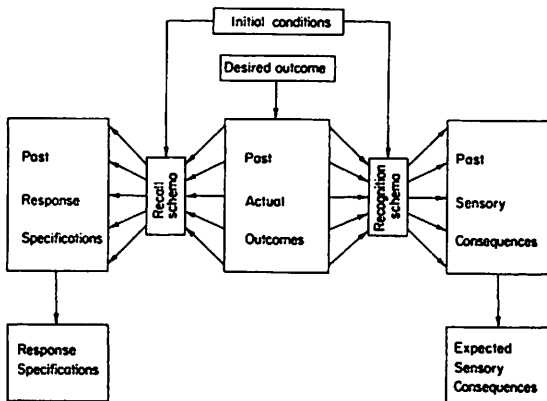


Figure 13. The recall and recognition schemas in relation to various sources of information. (From R.A. Schmidt, "The schema as a solution to some persistent problems in motor learning theory" in Motor Control: Issues and Trends, G.E. Stelmach (Ed.), 1976. Reprinted with permission from Academic Press, Inc. and the author.)

Analysis of visual input locates the target object within the subject's "reaching space". This location is fed to the control surface of the hand-reaching control system (i.e. it is not the job of this system to choose the target). On activation, the hand-reaching system directs a ballistic movement towards the target and activates a tuning mechanism to utilize visual and tactile feedback (cf. Figure 11).

(Figure 14)

But, prior to the actual reaching, analysis of visual input also extracts the size and orientation of the target object, and feeds them to the control surface of the grasping schema. This can be further decomposed into coactivation of finger and wrist adjustments, followed by inactivity until "contact" (i.e. when a portion of the object touches within the pre-shaped grasp) triggers the actual grasping movement, which shapes the hand on the basis of a subtle spatial pattern of tactile feedback.

The spoken instructions given to the subject drive the planning process that leads to the creation of the appropriate plan of action or coordinative structure -- which we here hypothesize to take the form of the distributed control program shown in the lower half of Figure 14, involving the interwoven activation of motor schemas for reaching and grasping. Note that the perceptual schemas hypothesized in the upper half of Figure 14 need not be regarded as a separate part of the coordinated control program. Rather, they provide the "identification algorithms" required to pass the proper parameter values to the motor schemas "per se". (For data consistent with the above general scheme, and which can be used to begin the refinement of Figure 14 into a testable model, see (83).)

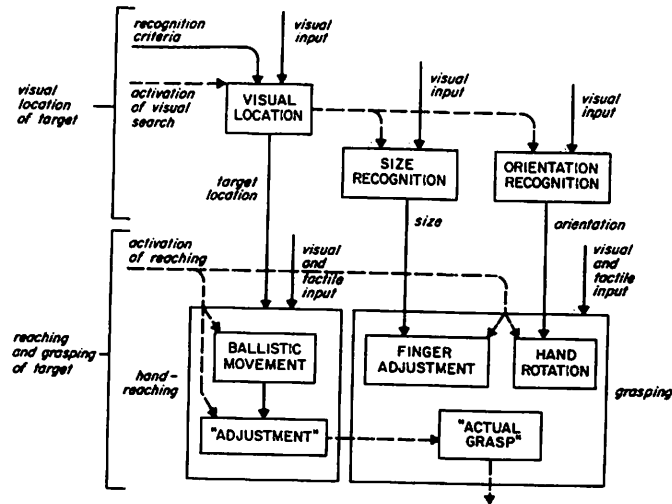


Figure 14. A hypothetical coordinated control program for a human's visually-directed reaching to grasp an object. - - -> control path; ———> data path.

FEEDFORWARD AND THE CEREBELLUM. Holmes (74) studied patients with one side of the cerebellum injured. Each patient was asked to stand in front of two vertical rows of lights and to move his hand swiftly back and forth. The normal subject, or the patient using the hand on the same side as the intact half of the cerebellum, exhibited a smooth trajectory of the kind shown in Figure 15A; while a patient using the hand whose side of the cerebellum had been grossly damaged exhibited the erratic trajectory of Figure 15B. One sees that the overall program for control of the trajectory is unimpaired by the cerebellar damage: the patient still moves from one target to the next in the required sequence. Instead, in the words of Holmes (74, p. 14):

"As a normal limb approaches its object its velocity declines at a uniform rate till it comes to rest, but the speed of the affected limb is often unchecked till the object is reached or even passed, and then the error is corrected by a series of jerky secondary movements, or the finger or toe may swing around the point to be touched without actually coming to rest."

Holmes attributes these effects to "irregular innervation of the prime movers", but I suggest that we might characterize the deficiency as a loss of the feedforward required to initiate active deceleration perhaps halfway through the movement to avoid overshooting. This transformation of spatial extent into proper timing of acceleration and deceleration was accomplished by the Braitenberg and Onesto (22) model of the cerebellum. Unfortunately, that model is not consistent with subsequent detailed studies of cerebellar structure and function, but it does remain a transformation whose mechanisms are to be sought in the cerebellum and its environs. Interestingly, a common strategy in these cerebellar patients seems to involve voluntary superposition of conscious feedforward control which is less accurate than the cerebellar mechanism it replaces:

"Frequently, however, the excursion of the limb is arrested too soon and the object is attained by similar correcting movements. This premature arrest is sometimes merely a manifestation of lack of uniformity in the rate of movement, but it seems to be commonly due to a

voluntary check: it occurs particularly in patients who have experienced the risk of overshooting the mark." (74, p. 15)

(Figure 15)




Figure 15. Execution of back and forth motion by a patient using hand on the same side as (A) the intact half of the cerebellum, (B) the half of the cerebellum badly damaged by a gunshot wound. (From G. Holmes, "The cerebellum of Man", Brain 62: 1-30, 1939. Reprinted with permission from Oxford University Press.)

WHEN IS A PROGRAM "HARD-WIRED"? Consider more carefully the mental processes required to complete the task of Figure 15:

- (1) The patient must comprehend the instructions (perhaps given by a combination of word and gesture) to establish (a) that a sequence of targets is to be pointed at in rapid succession; and (b) to define the targets visually in the correct order. This is the planning process to which we return in the section on artificial intelligence.
- (2) The execution of the task is accomplished in seven movements (where we refer to the six targets in order as d_1 through d_6):

- m_1 . Point at d_1 .
- \vdots
- m_6 . Point at d_6 .
- m_7 . Return hand to resting position.

As we have already suggested, each action in turn involves both a feedforward movement which in the normal subject yields approximately the desired result; and tuning of the pointing on the basis of visual and other feedback. Feedforward requires the use (at some stage prior to activation) of a visuo-motor map to convert visual coordinates into the pattern of muscular activity required to yield pointing in the given direction. Feedback requires a less accurate conversion of direction of error sensed visually to a direction of compensatory movement.

An interesting experiment would test the conjecture that another predictive component enters the motion beyond the feedforward already discussed:

that each movement is preshaped on the basis of what comes next (11). It is well-known that the utterance of a given phoneme is highly context-dependent. To get similar data on limb movement, we need to consider EMG and cine data on "Point at d_n " as a function of whether or not the subject is told to hold at d_n before moving again, and as a function of d_{n+1} .

Each step m_1 through m_6 is, we hypothesize, accomplished by a motor schema which uses both feedforward and feedback. In Figure 16, we show two different coordinated control programs for the sequential activation of the m_n 's. Figure 16A corresponds to the case in which the process of understanding the instructions leads to the creation of a program in which six explicit pointing activities are connected by an appropriate activation chain; while Figure 16B has a single pointing controller repeatedly given updated target information, and re-activated so long as another target remains. In view of our earlier discussion of "Maps as control surfaces", we may posit that each pointing is achieved by an increase of input activity at an appropriately target-related locus of some layered structure of the brain, and it may be moot as to whether we consider the first six control systems of Figure 16B as six separate systems, or as six subsystems (possibly even overlapping) of a single layered controller.

The patterns of activation-passing and data flow in Figures 14 and 16 are not to be thought of as "hard-wired" connections. Holmes' patients could just as well have been instructed to traverse the targets in a different order. In contrast to such "hard-wired" control systems as the α - γ system or the vestibulo-ocular reflex system or the frog's tectal controller of snapping, we have to posit a system which can construct a coordinated control program and use this coordinated control program to properly control the pattern of activation and message-passing of a whole

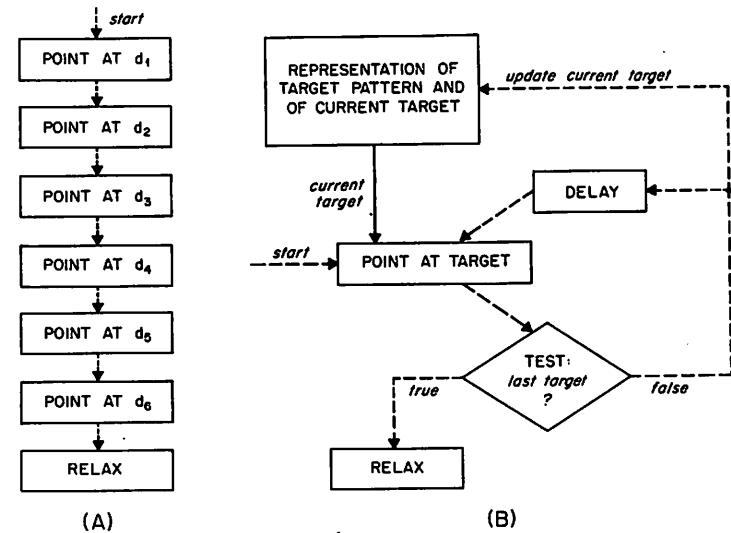


Figure 16. Two coordinated control programs for the pointing task of Figure 15.

variety of control systems. This planning and supervision is what seems unimpaired in the performance of Holmes' cerebellar patients, despite the gross impairment of the feedforward component of the motor schema for pointing.

(Figure 16)

In summary, then, a coordinated control program may be a representation of "hard-wired" interactions between neural subsystems; but for much human activity it represents a structure apart from the control system. It then controls the activation of structures elsewhere in the brain, possibly involving repeated activation of subsystems (as in an alternating movement task).

There are many different coordinated control programs compatible with a given (roughly described) motor behavior. The point, of course, is that such programs provide a language for the crisp presentation of alternatives. These must be put to experimental test. Unfortunately, the neural mechanisms for the "planned" coordinated control of motor schemas seem to be beyond the range of current experimental investigation. Below, we shall review certain AI approaches to planning which may provide a framework for the development of such investigations in the future. The separate structures created by planning may either be transient (e.g., created in response to verbal instructions for executing a seldom-repeated task) or may acquire a relatively permanent status being gradually refined through repeated execution (as in skill acquisition).

COMPETITION AND COOPERATION OF MOTOR SCHEMAS. A coordinated control program will, in general, involve both competition and cooperation of its constituent motor schemas. In Figures 14 and 16 we only see cooperation -- every system is activated in turn, and either controls overt behavior, or passes

necessary data to other systems. However, in a general situation (recall our discussion of "competitive computation of controllers"), the sensory input may activate many systems, and competition or planning processes will be required to determine which ones actually take part in overt behavior.

Skill Acquisition

Recalling our earlier discussion of "mutability of motor schemas", we briefly explore possible implications of the concept of coordinated control program for the analysis of skill acquisition.

The feedback provided to a subject learning a motor skill may be of little use when the overall structure of a coordinated control program has yet to be learnt no matter how appropriate it may be for fine tuning specific motor schemas within the program. Since we may posit that a system will learn best with feedback suited to its current learning strategy, study of the relative effectiveness of different types of feedback during skill acquisition may provide useful data for inferring what the strategy might be.

In a related study of skill acquisition (124), subjects were to learn to use a control stick to reproduce a curve on an oscilloscope. The curve was 4 seconds long, and its production involved 4 arm reversals. Training blocks were followed by tests. One group performed without feedback, and were then given a chance, after each block of trials, to compare the desired trajectory with that they had created. Three other groups received on-line feedback during each training block of trials -- one group had their arm driven through the pattern and were instructed to anticipate the movements; one group controlled the spot on the screen, but with a torque motor pulling the stick back on course if the error exceeded a threshold;

while the last group had visual feedback, as a tracking spot that migrated along the trajectory was shown on the scope. Because of the error-correction provided by such feedback, these three groups of subjects performed better than the first group during the acquisition trials. But during testing without such feedback, the subjects were found to learn the task no more quickly than the control subjects who essentially did the task open-loop. Pew commented that: "Subjects need information which they can use to improve their performance on subsequent trials, not information that contributes to better performance on the trial in progress." The feedback provided was not useful for the acquisition of the overall structure of the control program required to perform the task.

Such data are consistent with the hypothesis that a behavior may be under the control of a coordinated control program which determines how activation and data are to be passed between motor schemas. But we also posit that, within this overall control structure, the units may be mutable with parameters being updated over time by some neural analog of the control theorist's identification algorithm. Conscious control may enable us to "switch the overall program structure", but in many cases the tuning of the schemas must be "taken as is". The experimental delineation of the necessary neural mechanisms is a challenging one. A good start has been made on the mutability of the motor schemas themselves, but the study of the neural substructure of "overall program structure" has hardly begun.

Summary

We viewed feedforward both as a means whereby a controller can "pre-empt" the effect of disturbances, and as a strategy that generates large control signals that rapidly bridge large discrepancies in desired output. In discrete-activation feedforward, each activation of the feedforward controller

returns the controlled system to the "right ballpark" in which feedback can operate effectively. In co-activation feedforward, feedforward control is continually active to provide the approximate control signal required to maintain the desired output, with feedback serving to refine this approximation.

We then explored the notion of a coordinated control program in which control schemas are so scheduled that simultaneous actions are coordinated and successive actions are smoothly phased one into the next. While certain basic "programs" are "hard-wired" into the organism, most "programs" are generated as the result of an explicit planning process. A hypothetical program for reaching to a visually-located target showed that perceptual schemas need not necessarily be defined as separately specified as part of a coordinated control program, but may instead enter automatically as the identification algorithms required to define the appropriate values of the parameters in the motor schemas entering into the program.

The work of Holmes was cited to suggest that the overall structure of a coordinated control program may survive despite damage to (e.g. the feedforward components of) the neural circuitry that implements the motor schemas; and that conscious control may be able to counteract some of the resultant impairment.

While the neural mechanisms for the "planned" coordinated control of motor schemas seem to be beyond the range of current experimental investigation, we shall suggest below that AI approaches to planning may provide a framework for the development of such investigations in the future. Meanwhile, the notion of coordinated control program does make contact with the motor skills literature, the feedback being provided in learning being useful only if it matched the high-level structure of control program required to perform the task.

THE PERSPECTIVE OF ARTIFICIAL INTELLIGENCE

The progression of this chapter has been from concepts firmly rooted in neural experiments to concepts that appear, at this stage of our understanding, to be appropriate to functional analysis of "perceptual structures and distributed motor control" but not all of which are as yet amenable to experimental analysis. Our analysis of visuomotor coordination in frog and toad was based on a rich interaction of theory and experiment, while our analysis of perceptual and motor schemas could draw on a body of relevant experimental studies in neuroscience. When we turned to the notion of a coordinated control program, the link to experiments proved less well-developed, though we could see that the motor skills literature (137) could help us refine the functional analysis of such programs, even if little connection could be made with the neural level. In this last section, the citation of experiments is at most suggestive. We have seen that planning is a crucial ingredient in the control of movement; but have had to admit that little can be said about the neural mechanisms of planning (save, for example, that forebrain lesions impair it (94)). However, computer scientists have developed programs which, for example, enable a robot to plan its movement on the basis of visual information about its environment. To round out our presentation, then, it seems appropriate to briefly review some of the relevant literature. The discussion of experiments is far from definitive. Workers in AI (Artificial Intelligence) seldom interest themselves in neuroscience, and so there is very little literature on AI-neuroscience collaboration. We thus offer these concepts not in terms of successful neuroscience experiments, but rather to stimulate discussion of how to integrate planning notions into experiments on the neural control of movement.

Programs and Planners

COMPILERS AND INTERPRETERS. For certain motor tasks, a sufficiently high-level choice of strategy need not involve specification of the details of execution at the muscular level. For example, Lashley (89, p. 137) notes that animals which have learned a maze prior to injuries to the motor system continue to traverse it, although muscular activity may be changed completely -- one animal had to roll over completely in making each turn, yet made no errors in traversing the maze. This suggests that the brain can translate a high-level plan of the route to be traversed into different patterns of muscular activation as appropriate.

Computer scientists find it expedient to program a computer not in the machine language that directly controls the basic operations of its machinery but rather in terms of some high-level language. They provide the computer with a translation program which is either a "compiler" or "interpreter". A compiler is a program that allows the computer to translate a high-level program in its entirety prior to execution. An interpreter enables the computer to translate each instruction from the high-level program each time it is needed and then executes the resulting machine language subprogram. An intriguing question is to design experiments which could determine if the compiling vs. interpreting dichotomy is a useful one for describing neural processing of different skills.

PLANNING TECHNIQUES IN ARTIFICIAL INTELLIGENCE. In much computing, the process of translation is relatively straightforward -- multiplication of matrices, or sorting a sequence of numbers into a histogram can be carried out by standard procedures which (once the size of the matrices or the length of the sequence is given) do not depend on the idiosyncracies of the

situation. The study within AI of planning -- programming a computer to go from goal specifications to a plan for achieving the goals -- may be viewed as the study of "context-dependent translation". Consider, for example, the program do advance one step until door is reached as one way of translating the higher-level instruction go to the door. Such a translation is appropriate only if the system is facing the door and there are no obstacles en route. Otherwise, the position of the "organism" and the layout of obstacles within the room must be taken into account in determining a path.

We shall illustrate the AI approach to high-level planning of movement by considering a mobile robot equipped with a television camera and wheels. (A complementary area of AI is "scene analysis", the study of techniques for processing an image to build up an internal representation of the nature and location of objects in the scene of which the image is made (68).) The planning problem for such a robot is to enable it to respond to a command such as "Push the large cube to the door" by forming a plan that will take it to the cube without bumping into obstacles, and then will allow it to push the cube to the door without bumping into any obstacles. In compilation mode, planning will be completed prior to executing the plan. In interpretation mode, which is more in the spirit of the action/perception cycle, the plan will be continually updated as execution proceeds.

To simplify the planning problem for the mobile robot, we may schematize the world in which it finds itself by a graph in the sense of a set of points (called nodes) some of which are joined by directed edges. The nodes in this graph are targets for the movement of the robot, such as apertures and objects. The edges represent the transitions possible with available action choices. When given a task specified in a high-level language, the robot must first express this in terms of conditions upon a path through

the graph, such as to pass through certain nodes in a given order, while avoiding all nodes of another kind, and then it must find a path that satisfies these conditions and is also, ideally, as short as possible.

The simplest form of this problem is to find the shortest path from one specified node to another with no condition upon the intervening nodes. Doran and Michie (36) proposed a path-growing algorithm that used the idea of a heuristic distance, which is defined as an indication, rather than a guaranteed measure, of distance to the goal. For example, when one is travelling by road, the distance by air is a heuristic distance. The Doran-Michie approach is in the compilation mode: it proceeds by exploring alternative paths purely within its internal model (graph) of the environment, giving first preference to the exploration of paths whose next node has the least heuristic distance to the goal node. However, while this algorithm always yields a path from the start node (which represents where the system is) to the goal node (which represents where the system is to be), it cannot be guaranteed that the path obtained is the shortest. This problem was overcome in the A* algorithm (70), which gives first preference to the exploration of paths whose next node has a minimal sum of distance already traversed from the start node plus heuristic distance to the goal node.

In later work (46) a system called STRIPS was developed that, using theorem-proving techniques based on the representation of information about the environment in terms of predicate calculus, could take account of the more detailed constraints on a path that must be followed by a mobile robot. One sequel study (45) gave procedures for learning generalized robot plans. Changes in the relationship between the robot and the environment brought about by a sequence of actions were stored for each successful plan. Given sufficient data of this kind, various constants within both the sequence of

operations and the prescription of its effects could be replaced by variables to yield a number of MACROPS: high-level operations that could be used to simplify many familiar planning problems. A problem-solver called ABSTRIPS (132) uses planning in a hierarchy of abstraction spaces. It is usually more efficient to plan the overall stages of a series of movements, and then refine the plan with greater and greater detail, only making modifications when necessary, than to grow the overall plan step by step, with each step defined at the lowest level.

THE DIVERSITY OF PROGRAMS. With those programs written by AI researchers to explicitly embody aspects of the planning process, we are well on our way to an even more abstract type of program -- a program synthesizer, which takes as its input samples of behavior (e.g., input-output relations) and generates programs which can yield this type of behavior. In fact, we may regard the generation of MACROPS as an example of program synthesis. To summarize, then, we have the following levels of program sophistication:

- (i) Straight-line program: Executing a fixed predetermined sequence.
- (ii) Program: Executing a sequence of actions whose composition may depend on on-line testing of internal and external values.
- (iii) Planning = Context-dependent translation: Knowing what to do in general terms, find a solution tailored to the current situation.
- (iv) Program synthesis from examples: Starts by finding out what is the problem to be solved.

In adapting these concepts to the study of neural control of movement, we will certainly want to address specific questions such as varying speeds of performance, and postural adjustments. More importantly, we must learn how to transfer the above concepts from the serial style of most present-day computers to the distributed style of neural computation in which many

active subsystems are in continual interaction. But whether in computer or brain, we study programs which are part of the internal state of the system prior to action, and which can flexibly guide that action in terms of internal goals and external circumstances.

MAPPING THE ENVIRONMENT. Whereas localized activity in motor cortex seems related to activation of (small groups of) motor units, parietal lobe activity may be correlated with directed movements (106). In work that is still controversial, a number of workers have ascribed even more abstract forms of spatial representation to the hippocampus. Ranck (129) has shown that certain rat hippocampal cells will fire only when the animal is engaged in some specific activity, such as approaching its food dish or eating from it. O'Keefe and coworkers (113, 114, 115), on the other hand, find activity of rat hippocampal neurons correlated with the place at which the rat is in a maze. On this basis, O'Keefe suggests that the hippocampus holds a cognitive map representing the location of objects in the sort of non-egocentric coordinates one sees on a map in the usual geographic sense. A noteworthy feature of the studies of cognitive maps in the hippocampus is that it is relatively easy to correlate cellular activity in hippocampus with a location in the animal's current environment. One is thus tempted to posit that the hippocampus holds a "chart" of that local neighborhood -- which must, presumably, be read in from a whole "atlas" stored in long-term memory. Given Kandel's (85) data in *Aplysia* of short-term memory as a pattern of synaptic efficacy rather than of reverberation, it may be fruitful to design experiments which seek to show the synaptic settings of the current "chart" being established by patterned synaptic facilitation from elsewhere in the brain.

An important concept for the analysis of spatial mapping in the brain is that of "natural coordinates". In representing data concerning an animal's spatial behavior, it is tempting for the scientist to qualify them using the Cartesian coordinates of a Euclidian space. Yet care must be taken not to ascribe these coordinates to the neural representation -- for example, human behavior would often seem to be better described in terms of "Go to the door, from there you can see the fridge, in which you will find the piece of chicken", rather than "Advance 5.7 meters, turn through an angle of 72°, advance 4.2 meters, ...". From this viewpoint, we may urge experiments that seek to reconcile Ranck's and O'Keefe's findings by testing the hypothesis that the coordinates for hippocampal activity are egocentric and action-oriented -- so that places in the animal's world are not defined in terms of precise metrical relationships, but rather in terms of the behavior required to go from one place to another.

THE MONKEY AND THE BANANA. Another perspective on the question of natural coordinates comes from an AI "solution" to the "monkey and banana problem" made in the field of artificial intelligence. A monkey in a cage is considered to show "insight" when, in order to get a banana suspended from the ceiling, he pushes a box over and climbs upon it to reach the banana. Ernst and Newell (37) sought to show that a computer program called GPS was sufficiently flexible to solve this problem. (For a full description of GPS -- the General Problem Solver -- and its application to the analysis of human problem solving, see (110).) GPS is similar to the heuristic search problem above, in that problems can be expressed in terms of going from an initial state to some final state. To solve a particular class of problems, GPS needs three types of information: the differences between states that must be removed in solving the problem; the operators which

are at the system's disposal; and an operator-difference table which lists for each difference those operators that will often prove useful in removing it, though are by no means guaranteed to do so on every occasion. GPS can then use this information to construct, in a relatively efficient way, a sequence of operators whose application will indeed solve the problem. (It should be mentioned that in the last 15 years, AI has improved its repertoire of knowledge representations (13) and searching strategies (111), so that modern problem-solving programs are far more powerful than GPS.)

Unfortunately, the attempt to place the "monkey and banana" in the GPS framework trivializes the problem -- for, when provided with the data for this problem, GPS is given only two differences, vertical and horizontal, and told that "standing on the box" is the only operator which reduces vertical difference. This avoids the really interesting problems: What are the natural coordinates for the monkey? How does it come to learn that it can push the box around? How can it go from the knowledge that when on the box it will be further from the floor to the realization that the box is a tool for getting it closer to the ceiling? AI modelling designed to address these problems may well contribute to the development of the appropriate vocabulary for designing and evaluating experiments on the true coordinates of the posited hippocampal map and of other maps in the brain.

MAPS AND TOPOLOGIES. In elementary geography a map is a representation on paper of spatial relations of some part of the Earth's surface. However, maps are meaningless unless we have a process for using them. To get from Town C to Town D using a map, we must be able to:

1. recognize the representation of the towns,
2. choose a path between them, and
3. transform this into a plan for getting from C to D.

In a sense, then, the essence of a map is its path structure, and the utility of a map is that map-paths can be turned into programs for directing action.

The usual concept of a map in the brain simply correlates visual or motor targets with positions in a neural manifold. The above analysis suggests that we must demand more -- we must probe how the brain can encode sequences of action required to pass from one state to another, and look for neural circuitry which embodies a process for using these relationships. The distinction is like that which mathematicians make between a set -- a collection of unrelated elements -- and a topological space -- a set with an extra structure called a topology which expresses the connectedness relation between points of the set.

DEVELOPING A HIGH-LEVEL VOCABULARY. To briefly sketch progress towards a vocabulary for discussing the behaving organism's representation of the layout of the environment we consider the task of walking down a road. To follow such instructions as "Go north till you see the red barn," we need a map of elastic, rather than rigid, relationships. Moreover, we do not need a "full" map, for we can use the following strategy in unfamiliar territory. We scan buildings as we walk in anticipation of the recognition of a pattern of an expected class -- "red barn" -- rather than trying to analyze the entire scene in detail. The plan in this case specifies a sequence of landmarks. On each edge of the plan (to use the graph language of our discussion of AI planning techniques), we must use local guidance features: "Go north ..." requires (a) crude orientation to go the right way along the street; and (b) use of cues to stay on the footpath. During such traversal, the system must avoid collisions by computing a "free zone" -- a region into which a sudden movement may be made to avoid an obstacle

or collision. At any time, our current viewpoint offers only a perspective, and this visual information must be transformed to update our representation of the disposition of objects in the three-dimensional environment around us. Inverting the transformation, we have an expectation of visual input which both aids recognition of landmarks we seek and speeds recognition of those unexpected aspects of the environment which "demand" attention. For example, a speed-distance attention mechanism would ensure attention to a discrepancy (object not represented in the model) if the object's velocity is sufficient to carry it into our path. (Recall our earlier discussion of how time until contact may be computed from the optic flow.) In relating ego-motion to perception one should also note the distinction of two kinds of motion: locomotion and manipulation. Trevarthen (141) has suggested that the brain has two corresponding visual processes: one ambient, determining space at large around the body, the other focal which provides detail on small areas of space. But note that this corresponds to a subdivision of the current visual field, and does not involve the wider cognitive map, which involves representation of hidden objects. The "seen" and "unseen" worlds seem to involve different types of map. Arbib and Lieblisch (2), concerned with the motivational learning of spatial behavior, represented the cognitive map as a "world graph", with nodes corresponding to a recognizable place in the animal's world (but where a single place may be represented by more than one node) while each edge represents a path from one recognizable place to the next. The posited mechanisms for behavior were based on local weighting of nodes reachable from the current node. Kuipers (87) offered a multiple-frame representation for a system's spatial knowledge, with each "frame" (102) representing a perspective obtainable from a specific place while looking in a specific direction, together with information on how to get from one viewpoint to another.

Program Synthesis and Visuomotor Coordination

To conclude our somewhat speculative tour of AI concepts which may feed into the analysis of perceptual structures for motor control, we discuss how the computer scientist's notion of program synthesis might provide a setting for the study of development of visuomotor coordination (9).

AN EXPERIMENT ON VISUOMOTOR COORDINATION. In the set-up used by Held and Bauer (72) (Figure 17), a monkey must learn to grasp the handle, out of sight beneath the opaque shield, and pull it to bring the candy bar within licking distance. The initial position of the candy (with the handle beneath it) is the same at the start of each trial -- 45° from the midline. During training, the experimenter guides the monkey's hand to grasp and pull the handle. The monkey was constrained to use the same hand at each trial.

Animals were from two groups: experimental animals reared without ever seeing their hands, and control animals which were reared with their hands in view. Briefly, experimental animals were reluctant to reach; they reached in a groping or flailing manner; and they often grasped air on touching the handle; but they did perform quite well after successful grasping. (This reluctance to reach suggests that they monkeys were not only deprived of visuomotor coordination but were also deprived of tactilely-guided motor exploration. As we shall discuss below, a rich experience of tactile exploration might drastically change the animals' behavior. We need experiments which are closely coupled with theory in the analysis of the differential effects of the impact of various regimes of early experience.) By contrast, the control animals soon performed as well as normal animals





Figure 17. Top: Monkey reaching open-loop for manipulandum while gazing at target-reinforcer. Bottom: Monkey pulling manipulandum so as to bring target-reinforcer to his mouth. (From R. Held and J.A. Bauer "Development of sensorially-guided reaching in infant monkeys", *Brain Research* 71: 265-271, 1974. Reprinted with permission from Elsevier/North-Holland Biomedical Press BV and the authors.)

who could see their hands during the experiment. They would reach without hesitation and would open their hand in anticipation of contact with the handle.

PROGRAM SYNTHESIS AND THE AVAILABLE REPERTOIRE. In what follows we shall use three concepts: the notion of an adaptable motor schema; the computer scientist's notion of program synthesis (going from the examples provided by the experimenter's guidance of the monkey's hand to a program in the available repertoire of instructions); and the notion of a map as a control surface. The style of this subsection will be hypothetical, aimed at encouraging the assimilation of AI concepts into the study of motor control.

Our first hypothesis, then, is that in the experimental setting (72) all monkeys quickly learn that the following program works:

1. Grasp the handle.
2. Pull it bodyward.

We hypothesize that both groups have these instructions -- grasp and pull -- in their repertoire, and can use them for any hand position. The real difference in the groups comes in matching the precondition for instruction 1 -- getting the hand in position to grasp the handle.

As Held and Bauer observe, the data show that the control animal can match arm direction to visual direction, while the experimental animal cannot, and we represent this in the language of programs by saying that the control animal can expand 1 to:

- 1A. Reach below the candy.
- 1B. Grasp the handle.

Note that 1A is defined by a hand-eye coordination, and is meaningless to experimental monkeys -- for them, "below the candy" has meaning for eye movements, but cannot be used as a variable for an arm command. Instead,

the experimental animal is conditioned to approximate the rewarded arm direction. It thus builds up the ability to do a new command called "A-reach". (I use the neutral term "A-reach" rather than "reach in a direction of about 45°" to stress that we cannot assume the animal has an "angular map" to relate 45° to the command.) Note that an animal without visuomotor coordination but with a rich experience of reaching for objects in the dark would, presumably, not require such conditioning to acquire the A-reach, but could produce it from its prior repertoire. The current training would simply get the animal to learn which well-trained movement was appropriate.

Since the A-reaching of the experimental animal is inaccurate, we hypothesize that it is embedded in the following:

- 1'A. A-reach.
- 1'B. do move back and forth until grasp handle.

Note, then, that the control animal can use a command from its available repertoire, but the experimental animal must synthesize a new command.

Held and Bauer report a number of further experiments -- and in each case, the results sit comfortably in the framework of the above discussion.

(i) Spatial Generalization: After the initial learning, the animals were tested with seven positions of the candy bar and handle, the positions being 10° apart. The controls had little difficulty reaching directly for the handle in any position -- as one would expect if they were following the program [1A; 1B, 2] -- except (for reasons I do not understand) for the position nearest the midline.

The experimental monkeys reached toward the original position, 45° from the midline, and then swept the arm back and forth until it touched the handle -- as one would expect if they were following the program

[1'A; 1'B; 2] -- which worked quite well for target positions within 15° of the original.

(ii) Intermanual Transfer: With the candy and handle set up on the other side of the midline, it took several day's training to get all monkeys to use the other arm. However, once the change of arm was effected, control animals showed a high degree of intermanual transfer to all directions and were virtually errorless. On the other hand, experimental animals showed a marked decrement in performance, and the same relative inaccuracies as they had with the initially trained hand. This is as one would expect with the programs above.

With continued experience with targets in various directions, the experimental animals continued to improve their reaching to targets in varied directions -- but by slow and painstaking conditioning with extinction if testing was discontinued for a few weeks. In our programming terminology, we may say that the experimental animal has constructed a set of isolated commands, including the "A-reach" (our neutral terminology for the first-learned reaching at an angle of about 45°):

A-reach

B-reach

C-reach, etc.,

each paired with a different visual stimulus (or head direction for looking directly at the candy). By contrast, we posit that the control animal has a single command:

reach (X)

with a parameter, the visual coordinate X of the target.

FROM SET TO TOPOLOGICAL SPACE. With a few days of free sight of its limbs, an experimental animal can perform almost as well as a control animal. What has changed to make the isolated points of the original pattern cohere into a "map" which makes the visual direction parameter X available to an integrated reaching command?

One hypothesis is that the animal builds up "a map of sets", laboriously constructed point by point, i.e. that all that has changed is that the animal can reach for more points and more often than was possible in the experimental situation. However, I would hypothesize that the crucial difference is that the animal builds up a "path structure". Error-correction of a movement is based on a network of nearness relations. In the normal monkey there is a genuine visuomotor map of reaching; whereas (while it is denied visually-guided reaching) in the experimental monkey this is not so. Its visual system may contain a map for eye movements; and the monkey can certainly reach in different directions. However, the set of reachings has not been "topologized", since the animal has no experience of "directed" arm movements save those that bring the hand to, and move it about, the body surface. Our second hypothesis, then, stresses the importance of continuity and feedback in building up the map -- in this case, in making the visual parameter X available to the reaching command. This requires building up the necessary path relationships -- correlating limb movement with visual change, and thus supplying error signals to a hand-control system.

New hypotheses are fruitful, *inter alia*, if they lead to new experiments. Here, then, are two experiments. The first is neutral with respect to the "set-construction" vs. "path-space" hypotheses; whereas the second experiment would fail if only "set-construction" held.

(i) In the first set-up, have the opaque shield so modified that a light goes on atop the shield above the current position of the monkey's hand. Although it does not see its hand, the monkey should learn how its hand movements affect the movement of lights on the table. It would thus be able to synthesize step 1 of the program as

1"A. Move hand to turn on light nearest candy bar.

1"B. Grasp handle.

and would probably be defective only in the dexterity and anticipation of its grasping. I predict that its performance on spatial generalization and intermanual transfer would match that of the control animal.

(ii) The second experiment is more subtle, and is based on giving the animal a "tactile world" beneath the shield. Let the animal induce a "topology" by having an array of tactile targets beneath the shield, perhaps with texture gradients, and with rewards at various places. I predict that such an animal would perform moderately well on the first target position -- reaching for the handle at tactile position X -- and would probably suffer a slight decrement at the second position, since the change in visual position of the candy provides no cue as yet of a change in tactile position. However, I predict that his performance would improve dramatically after reaching criterion on several (vision, touch) coordinate pairs, for he would then have enough data to use the two topologies to extrapolate the correspondence, at least approximately, over a large part of the continuum.

This still leaves the question of why the monkey deprived of hand-vision cannot use proprioceptive feedback to "induce the topology" of arm positions. One answer might be that, given the apparent lack of motor

tasks for the experimental animal's limbs, no map has been established between proprioceptive feedback and motor reaching. Given a tactile world, the proprioceptive cues might well become available as part of the non-visual motor topology.

Summary

While some attention was given to behavioral and physiological analysis of cognitive mapping and visuomotor coordination, the main thrust of this section was to present a number of concepts from computer science and artificial intelligence (AI) which have been exposed to little in the way of neuroscientific experiment but which may serve to stimulate new models of perceptual structures and their role in the planning and control of movement.

We first built on the earlier discussion of the fact that "programs need not be stereotypes", and discussed three types of high-level "program-generating programs".

- 1) Compilers and interpreters afford two styles of translation programs, enabling a computer to perform a computation in terms of basic low-level steps on the basis of high-level program.
- 2) Planners are context-dependent translators which take an explicit statement of goals and a model of the environment to yield a program of action which will achieve the goals within the environment so represented. Most AI planners are like compilers, completing the entire plan prior to its execution. For modelling the action/perception cycle, we would seem to need planners that operate in the interpretive mode, with the plan being continually updated as action enables the organism to update its model of the environment.

- 3) A program synthesizer is a program designed to take samples of input-output behavior, and produce a program which will perform in the manner so exemplified. A program synthesizer is akin, in some ways, to the identification algorithm of an adaptive controller, but has far greater flexibility than the setting of parameters within some prespecified control scheme.

We then turned to the problem of "cognitive mapping" stressing that the "natural coordinates" in such a map will often look far more like programs for getting from one point to another than like the Euclidean coordinates the experimenter might use in quantifying behavioral data. From this perspective, we saw that the acquisition of visuomotor coordination might be analyzed in terms of both program synthesis and the learning of new motor schemas, rather than simply in terms of point-by-point pairing of layers of neurons.

CONCLUSION

We have offered a perspective on the role of visual information in the control of movement from the perspective of brain theory, seeking to bridge the gap between studies of behavior and overall function (motor skills, artificial intelligence and cognitive psychology) and the study of physiologically and anatomically well-defined neural nets (biological control theory and neural modelling).

We took biological control theory as our point of departure. The notion of state in control theory serves as antidote to the Sherringtonian analysis of the organism in reflex terms. However, we studied visuomotor coordination in frog and toad to stress the need to move beyond the lumped input of the systems studied in much biological control theory. Rather, we stress that visual input has rich spatial and temporal structure, and that specific processes are required to extract information relevant to the control of different behaviors. We thus see the many visual systems providing input to the control surfaces of a variety of controllers of different behaviors. The computation within each system can be analyzed in terms of competition and cooperation of neural nets; the complementary problem is that of how these different controllers are coordinated in determining the integrated activity of the organism.

Complexity seems to demand intermediate levels of analysis to mediate the flexibility of an animal's behavior in a richly structured environment. We saw that the environment could be represented in terms of an assemblage of "instantiations of perceptual schemas", each representing relevant parameters of some particular object or domain of interaction relevant to the organism. We used "planning" to refer to the process whereby a plan

of action is built up on the basis of the organism's goals and current perceptions; stressing that as the animal moves it updates the schema-assemblage that represents its environment, leading in turn to an updated plan, in a continual action/perception cycle. We used the term "motor schema" to denote the units from which each plan of action is constituted, and introduced the notion of "distributed control program" to suggest how indeed the motor schemas could be interwoven. We marshalled data to suggest the utility of the control-theoretic notions of "feedforward" and "identification algorithm" in describing the structure of motor schemas at a level that can connect with neurophysiological investigation. The growing subtlety of experiments in the motor skill literature, especially when coupled with EMG recording, begin to let us essay functional accounts of "coordinated control programs" for a number of behaviors that involve the integration of several motor schemas. But the neural implementation of such programs -- as distinct from the motor schemas themselves -- remains elusive, and we concluded with an account of planning and the acquisition of cognitive maps in the vocabulary of Artificial Intelligence to suggest one possible direction for the development of the proper vocabulary for the design of experiments on those systems that are not "hard-wired".

We thus see a spectrum of concepts from those firmly rooted in the findings of the neuroscientist's laboratory to those which have proved their efficacy in the design of robots but whose relevance to the analysis of brain remains speculative. The aim of this chapter has been to encourage the interaction of brain theorist and experimentalist in evolving and adapting these latter concepts to let us further understand the neural basis for the role of perceptual structures in the distributed control of movement.

REFERENCES

1. ALLUM, J. H. J. Responses to load disturbances in human shoulder muscles: the hypothesis that one component is a pulse test information signal. Exp. Brain Res. 22: 307-326, 1975.
2. AMARI, S. AND M. A. ARBIB. Competition and cooperation in neural nets. In: Systems Neuroscience, edited by J. Metzler. New York: Academic Press, 1977, p. 119-165.
3. APTER, J. The projection of the retina on the superior colliculus of cats. J. Neurophysiol. 8: 123-134, 1945.
4. APTER, J. Eye movements following strychninization of the superior colliculus of cats. J. Neurophysiol. 9: 73-85, 1946.
5. ARBIB, M. A. Transformation and somatotopy in perceiving systems. Proc. Second Int. Joint Conf. Artificial Intelligence, 1971, p. 140-147.
6. ARBIB, M. A. The Metaphorical Brain: An Introduction to Cybernetics as Artificial Intelligence and Brain Theory. New York: Interscience, 1972.
7. ARBIB, M. A. Artificial intelligence and brain theory: unities and diversities. Ann. Biomed. Eng. 3: 238-274, 1975.
8. ARBIB, M. A. From automata theory to brain theory. Int. J. Man-Mach. Stud. 7: 279-295, 1975.
9. ARBIB, M. A. Program synthesis and sensorimotor coordination. Brain Theory Newsletter 2: 31-33, 1976.

10. ARBIB, M. A., C. C. BOYLLS AND P. DEV. Neural models of spatial perception and the control of movement. In: Cybernetics and Bionics, edited by W. D. Keidel, W. Händler and M. Spreng. München/Wein: Oldenbourg, 1974, p. 216-231.
11. ARBIB, M. A., G. F. FRANKLIN AND N. NILSSON. Some ideas on information processing in the cerebellum. In: Neural Networks, edited by E. R. Caianiello. Berlin/New York: Springer-Verlag, 1968, p. 43-58.
12. ARBIB, M. A. AND I. LIEBLICH. Motivational learning of spatial behavior. In: Systems Neuroscience, edited by J. Metzler. New York: Academic Press, 1977, p. 221-239.
13. ASHBY, W. R. Homeostasis. In: The Encyclopaedia of the Biological Sciences, edited by P. Gray. Reinhold, 1961, p. 487, 488.
14. BARLOW, H. B., C. BLAKEMORE AND J. D. PETTIGREW. The neural mechanism of binocular depth discrimination. J. Physiol. 193: 327-342, 1967.
15. BARTLETT, F. C. Remembering. London/New York: Cambridge University Press, 1932.
16. BERNARD, C. Leçons sur les Phénomènes de la Vie. Balliere, 1878.
17. BERNSTEIN, N. A. The Coordination and Regulation of Movements. (Translated from Russian), Pergamon, 1967.
18. BOBROW, D. G. AND A. COLLINS. Representation and Understanding: Studies in Cognitive Science. New York: Academic Press, 1975.

19. BOYLLS, C. C. A theory of cerebellar function with applications to locomotion. I. The physiological role of climbing fiber inputs in anterior lobe operation. COINS Tech. Report 75C-6, Univ. of Massachusetts at Amherst, 1975.
20. BOYLLS, C. C. A theory of cerebellar function with applications to locomotion. II. The relation of anterior lobe climbing fiber function to locomotor behavior in the cat. COINS Tech. Report 76-1, Univ. of Massachusetts at Amherst, 1976.
21. BOYLLS, C. C. Prolonged alterations of muscle activity induced in locomoting preamillary cats by microstimulation of the inferior olive. Brain Res. (in press).
22. BRAITENBERG, V. AND N. ONESTO. The cerebellar cortex as a timing organ. In: Proc. Congresso di Inst. di Medicina Cibernetica, Naples, 1969, p. 239-255.
23. BROOKS, V. B. Some examples of programmed limb movements. Brain Res. 71: 299-308, 1974.
24. CANNON, W. B. The Wisdom of the Body. Norton, 1939.
25. COLLETT, T. Stereopsis in toads. Nature 267: 349-351, 1977.
26. CRAIK, K. J. W. The Nature of Explanation. Cambridge Univ. Press, 1943.
27. CRAIK, K. J. W. Theory of the human operator in control systems. I. The operator as an engineering system. Br. J. Psychol. 38: 56-61, 1947.
28. CRAIK, K. J. W. Theory of the human operator in control systems. II. Man as an element in a control system. Br. J. Psychol. 38: 142-148, 1948.

29. CRAIK, K. J. W., AND M. A. VINCE. Psychological and physiological aspects of control mechanisms with special reference to tank gunnery. Part I. UK MRC Military Personnel Res. Comm. Report, London, England, 1943. Ergonomics 6: 1-33, 1963.
30. CRAIK, K. J. W. AND M. A. VINCE. Psychological and physiological aspects of control mechanisms. Part II. UK MRC Military Personnel Res. Comm. Report BPC 44/322, London, England, 1944. Ergonomics 6: 419-440, 1963.
31. CREED, R. S., D. DENNY-BROWN, J. C. ECCLES, E. G. T. LIDDELL AND C. S. SHERRINGTON. Reflex Activity of the Spinal Cord. Oxford Univ. Press, 1932 (reprinted with annotations by D. P. C. Lloyd in 1972).
32. DEV, P. Computer simulation of a dynamic visual perception model. Int. J. Man-Mach. Stud. 7: 511-528, 1975.
33. DIDDAY, R. L. The Simulation and Modelling of Distributed Information Processing in the Frog Visual System (Ph.D. Thesis). Stanford Univ., 1970.
34. DIDDAY, R. L. A model of visuomotor mechanisms in the frog optic tectum. Math. Biosci. 30: 169-180, 1976.
35. DIDDAY, R. L. AND M. A. ARBIB. Eye-movements and visual perception: a "two-visual system" model. Int. J. Man-Mach. Stud. 7: 547-569, 1975.
36. DORAN, J. AND D. MICHIE. Experiments with the graph traverser program. Proc. Royal Soc. London, Series A 294: 235-259, 1966.
37. ERNST, G. W. AND A. NEWELL. GPS: A Case Study in Generality and Problem Solving. New York: Academic Press, 1969.

38. EWERT, J.-P. Untersuchungen über die Anteile zentralnervöser Aktionen an der taxisspezifischen Ermüdung beim Beutefang der Erdkröte (Bufo Bufo L.). Z. Vergl. Physiol. 57: 263-298, 1967.
39. EWERT, J.-P. Zentralnervöse Analyse und Verarbeitung visueller Sinnesreize. Naturwiss. Rundsch. 25: 1-11, 1972.
40. EWERT, J.-P. The visual system of the toad: Behavioral and physiological studies on a pattern recognition system. In: The Amphibian Visual System, A Multidisciplinary Approach, edited by K. V. Fite. New York: Academic Press, 1976, p. 141-202.
41. EWERT, J.-P. AND W. VON SEELEN. Neurobiologie und System-Theorie eines visuellen Muster-Erkennungsmechanismus bei Kröten. Kybernetik 14: 167-183, 1974.
42. EWERT, J.-P. AND A. VON WIETERSCHEIM. Musterauswertung durch Tectum- und Thalamus/Praetectum-Neurone im visuellen System der Kröte (Bufo Bufo L.). J. Comp. Physiol. 92: 131-148, 1974.
43. EWERT, J.-P. AND A. VON WIETERSCHEIM. Der Einfluss von Thalamus/Praetectum-Detekten auf die Antwort von Tectum-Neuronen gegenüber visuellen Mustern bei der Kröte (Bufo Bufo L.). J. Comp. Physiol. 92: 149-160, 1974.
44. FEL'DMAN, A. G. Functional tuning of the nervous system with control of movement or maintenance of a steady posture -- II. Controllable parameters of the muscles. Biophysics 11: 565-578, 1966 (Translated from Biofizika 11: 498-508, 1966).

45. PIKES, R. E., P. E. HART AND N. J. NILSSON. Learning and executing generalized robot plans. Artif. Intell. 3: 251-288, 1972.
46. PIKES, R. E. AND N. J. NILSSON. STRIPS: A new approach to the application of theorem proving to problem solving. Artif. Intell. 2: 189-208, 1971.
47. FITCH, H. L. AND M. T. TURVEY. On the control of activity: some remarks from an ecological point of view. In: Psychology of Motor Behavior and Sport - 1977, edited by D. M. Landers and R. W. Christina. Urbana, Ill.: Human Kinetics Publishers, 1978, p. 3-35.
48. FREDERIKS, J. A. M. Disorders of the body schema. In: Handbook of Clinical Neurology. Disorders of Speech, Perception, and Symbolic Behaviour. North-Holland Publishing Co., 1969, vol. 4, p. 207-240.
49. GEL'FAND, I. M., V. S. GURFINKEL', M. L. SHIK AND M. L. TSETLIN. Certain problems in the investigation of movement. In: Automata Theory and Modeling of Biological Systems. Academic Press, 1973, p. 160-171. (Translated from the Russian original Models of the Structural-Functional Organization of Certain Biological Systems. Nauka, 1966.)
50. GIBSON, J. J. The optical expansion-pattern in aerial location. Amer. J. Psychol. 68: 480-484, 1955.
51. GIBSON, J. J. Visually controlled locomotion and visual orientation in animals. Br. J. Psychol. 49: 182-194, 1958.
52. GIBSON, J. J. The Senses Considered as Perceptual Systems. London: George Allen & Unwin, 1966.

53. GIBSON, J. J. The theory of affordances. In: Perceiving, Acting and Knowing: Toward an Ecological Psychology, edited by R. E. Shaw and J. Bransford. Hillsdale, N.J.: Erlbaum, 1977.
54. GIBSON, J. J. AND E. J. GIBSON. Continuous perspective transformation and the perception of rigid motion. J. Exp. Psychol. 54: 129-138, 1957.
55. GILBERT, P. How the cerebellum could memorise movements. Nature 254: 688-689, 1975.
56. GILBERT, P. F. C. AND W. T. THACH. Purkinje cell activity during learning. Brain Res. 128: 309-328, 1977.
57. GORDON, D. A. Static and dynamic visual fields in human space perception. J. Opt. Soc. Am. 55: 1296-1303, 1965.
58. GREENE, P. H. Seeking mathematical models of skilled actions. In: Biomechanics, edited by R. Brodsky. Plenum Press, 1969, p. 149-180.
59. GREENE, P. H. Introduction. In: Models of the Structural-Functional Organization of Certain Biological Systems, edited by I. M. Gel'fand, V. S. Gurfinkel, S. V. Fomin and M. S. Tsetlin. MIT Press, 1971, p. xi-xxxi. (Translated from the Russian by C. R. Beard.)
60. GREENE, P. H. Problems of organization of motor systems. In: Progress in Theoretical Biology, edited by R. Rosen and F. M. Snell. Academic Press, 1972, vol. 2, p. 303-338.
61. GREGORY, R. L. On how so little information controls so much behavior. In: Towards a Theoretical Biology, 2: Sketches, edited by C. H. Waddington. Edinburgh: Edinburgh Univ. Press, 1969.

62. GRILLNER, S. Locomotion in vertebrates: central mechanisms and reflex interaction. Physiol. Rev. 55: 247-304, 1975.
63. GROSSBERG, S. On learning of spatiotemporal patterns by networks with ordered sensory and motor components, I. Excitatory components of the cerebellum. Stud. Appl. Math. 48: 105-132, 1969.
64. GROSSBERG, S. Neural expectation: cerebellar and retinal analogs of cells fired by learnable or unlearned pattern classes. Kybernetik 10: 49-57, 1972.
65. GROSSBERG, S. Adaptive pattern classification and universal recoding, I. Parallel development and coding of neural feature detectors. Biol. Cybernetics 23: 121-134, 1976.
66. GROSSBERG, S. Adaptive pattern classification and universal recoding, II. Feedback, expectation, olfaction, and illusions. Biol. Cybernetics 23: 187-202, 1976.
67. GRÜSSER, O.-J. AND U. GRÜSSER-CORNEHLS. Physiology of the anuran visual system. In: Neurobiology of the Frog, edited by R. Llinás and W. Precht. Springer-Verlag, 1976.
68. HANSON, A. R. AND E. M. RISEMAN. Computer Vision Systems. New York: Academic Press, 1978.
69. HARKNESS, L. Chameleons use accommodation cues to judge distance. Nature 267: 346-349, 1977.
70. HART, P., N. J. NILSSON AND B. RAPHAEL. A formal basis for the heuristic determination of minimum cost paths. IEEE Trans. Syst. Sci. Cybern. SSC-4: 100-107, 1968.

71. HEAD, H. AND G. HOLMES. Sensory disturbances from cerebral lesions. Brain 34: 102-254, 1911.
72. HELD, R. AND J. BAUER. Development of sensorially-guided reaching in monkeys. Brain Res. 71: 265-271, 1974.
73. HOLLERBACH, J. M. The minimum energy movement for a spring muscle model. AI Memo. 424, Cambridge, Mass.: MIT A.I. Lab., 1977.
74. HOLMES, G. The cerebellum of Man. Brain 62: 1-30, 1939.
75. INGLE, D. Visual releasers of prey-catching behavior in frogs and toads. Brain Behav. Evol. 1: 500-518, 1968.
76. INGLE, D. Spatial vision in anurans. In: The Amphibian Visual System, edited by K. V. Fite. New York: Academic Press, 1976, p. 119-141.
77. INGLE, D. Focal facilitation effects in visually elicited feeding by the frog. (In press).
78. INGLE, D. A functional approach to the many visual systems dilemma. (In press).
79. ITO, M. Neurophysiological aspects of the cerebellar motor control system. Int. J. Neurol. 7: 162-176, 1970.
80. ITO, M. Neural design of the cerebellar motor control system. Brain Res. 40: 81-84, 1972.
81. ITO, M. The control mechanisms of cerebellar motor systems. In: The Neurosciences: Third Study Program, edited by F. O. Schmitt and F. G. Worden. Cambridge, Mass.: MIT Press, 1974, p. 293-303.

82. ITO, M. Recent advances in cerebellar physiology and pathology. Ad. Neurol. 21: 59-84, 1978.
83. JEANNEROD, M. AND B. BIGUER. Visuomotor mechanisms in reaching within extrapersonal space. In: Advances in the Analysis of Visual Behavior, edited by D. J. Ingle, R. J. W. Mansfield and M. A. Goodale. Cambridge, Mass.: MIT Press (in press).
84. JULESZ, B. Foundations of Cyclopean Perception. Chicago: University of Chicago Press, 1971.
85. KANDEL, E. R. A Cell-Biological Approach to Learning. Bethesda, Md.: The Society for Neuroscience, 1978.
86. KILMER, W. L., W. S. McCULLOCH AND J. BLUM. A model of the vertebrate central command system. Int. J. Man-Mach. Stud. 1: 279-309, 1969.
87. KUIPERS, B. J. Representing knowledge of large-scale space. AI-TR-418, The AI Laboratory, MIT, 1977.
88. LAMETTRIE, J. Man a Machine. Chicago, 1953. (Trans. by G. Bussey from the 1748 French original.)
89. LASHLEY, K. S. Brain Mechanisms and Intelligence. Chicago: Univ. of Chicago Press, 1929.
90. LASHLEY, K. S. The problem of serial order in behavior. In: Cerebral Mechanisms in Behavior, edited by L. Jeffress. New York: Interscience, 1951, p. 112-136.
91. LEE, D. N. Visual information during locomotion. In: Perception: Essays in Honor of J. J. Gibson, edited by R. B. McLeod and H. L. Pick, Jr. Ithaca, N.Y.: Cornell Univ. Press, 1974, p. 250-267.

92. LEE, D. N. AND J. R. LISHMAN. Visual control of locomotion. Scand. J. Psychol. 18: 224-230, 1977.
93. LETTVIN, J. Y., H. MATURANA, W. S. McCULLOCH AND W. H. PITTS. What the frog's eye tells the frog's brain. Proc. IRE 47: 1940-1951, 1959.
94. LURIA, A. R. The Working Brain, An Introduction to Neuropsychology. (Translated from the Russian by B. Haigh), Penguin Books, 1973.
95. MacKAY, D. M. Cerebral organization and the conscious control of action. In: Brain and Conscious Experience, edited by J. C. Eccles. Springer-Verlag, 1966, p. 422-440.
96. MARR, D. A theory of cerebellar cortex. J. Physiol. 202: 437-470, 1969.
97. MARR, D. AND T. POGGIO. Cooperative computation of stereo disparity. Science 194: 283-287, 1977.
98. MAXWELL, J. C. On governors. Proc. R. Soc. London 16: 270-283, 1868.
99. McCULLOCH, W. S. AND W. H. PITTS. A logical calculus of the ideas immanent in nervous activity. Bull. Math. Biophys. 5: 115-133, 1943.
100. MILLER, G. A., E. GALANTER AND K. H. PRIBRAM. Plans and the Structure of Behavior. New York: Holt, Rinehart and Winston, 1960.
101. MINSKY, M. L. Steps towards artificial intelligence. Proc. IRE 49: 8-30, 1961.
102. MINSKY, M. L. A framework for representing knowledge, In: The Psychology of Computer Vision, edited by P. H. Winston. New York: McGraw-Hill, 1975, p. 211-277.

103. MONTALVO, F. S. Consensus vs. competition in neural networks. Int. J. Man-Mach. Stud. 7: 333-346, 1975.
104. MOUNTCASTLE, V. B. The world around us: neural command functions for selective attention. Neurosci. Res. Program Bull. 14: suppl. 1-47, 1976.
105. MOUNTCASTLE, V. B. An organizing principle for cerebral function: the unit module and the distributed system. In: The Mindful Brain, by G. M. Edleman and V. B. Mountcastle. Cambridge, Mass.: MIT Press, 1978.
106. MOUNTCASTLE, V. B., J. C. LYNCH, A. GEORGOPOULOS, H. SAKATA AND C. ACUNA. Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. J. Neurophysiol. 38: 871-908, 1975.
107. NASHNER, L. M. Analysis of stance posture in humans. In: Handbook of Behavioral Neurobiology. Plenum Press (in press).
108. NEISSER, U. Cognition and Reality: Principles and Implications of Cognitive Psychology. San Francisco: W. H. Freeman, 1976.
109. NELSON, J. I. Globality and stereoscopic fusion in binocular vision. J. Theoretical Biol. 49: 1, 1975.
110. NEWELL, A. AND H. A. SIMON. Human Problem Solving. Englewood Cliffs, N.J.: Prentice-Hall, 1972.
111. NILSSON, N. J. Problem-Solving Methods in Artificial Intelligence. McGraw-Hill, 1971.

112. NOTON, D. AND L. STARK. Scanpaths in eye movements during pattern perception. Science 171: 308-310, 1971.
113. O'KEEFE, J. Place units in the hippocampus of the freely moving rat. Exp. Neurol. 51: 78-109, 1976.
114. O'KEEFE, J. AND J. DOSTROVSKY. The hippocampus as a spatial map. Brain Res. 34: 171-175, 1971.
115. O'KEEFE, J. AND L. NADEL. The Hippocampus as a Cognitive Map. Oxford University Press, 1978.
116. OLDFIELD, R. C. AND O. L. ZANGWILL. Head's concept of the body schema and its application in contemporary British psychology. Br. J. Psychol. 32: 267-286, 1942; 33: 58-64, 113-129, 143-149, 1943.
117. OSCARSSON, O. The sagittal organization of the cerebellar anterior lobe as revealed by the projection patterns of the climbing fiber system. In: Neurobiology of Cerebellar Evolution and Development, edited by R. Llinás. American Medical Association, 1969, p. 525-537.
118. PAILLARD, J. AND M. BROUCHON. A proprioceptive contribution to the spatial encoding of position cues for ballistic movements. Brain Res. 71: 273-284, 1974.
119. PAL'TSEV, Y. I. Functional reorganization of the interaction of the spinal structure in connexion with the execution of voluntary movement. Biophysics 12: 313-322, 1967. (Translated from Biofizika 12: 277-284.
120. PELLIONISZ, A. Proposal for shaping the dynamism of Purkinje cells by climbing fiber activation (with a discussion by C. C. Boylls). Brain Theory Newsletter 2: 2-6, 1976.

121. PELLIONISZ, A. AND R. LLINÁS. A computer model of cerebellar Purkinje cells. Neuroscience 2: 37-48, 1977.
122. PELLIONISZ, A., R. LLINÁS AND D. H. PERKEL. A computer model of the cerebellar cortex of the frog. Neuroscience 2: 19-36, 1977.
123. PERENIN, M. T. AND M. JEANNEROD. Visual function within the hemianopic field following early cerebral hemidecortication in man. I. Spatial localization. Neuropsychologia 16: 1-13, 1978.
123. PETTIGREW, J. D., T. NIKARA AND P. O. BISHOP. Binocular interaction on single units in cat striate cortex. Exp. Brain Res. 6: 391-410.
124. FEW, R. W. Levels of analysis in motor control. Brain Res. 71: 393-400, 1974.
125. PIAGET, J. Biology and Knowledge: An Essay on the Relations between Organic Regulations and Cognitive Processes. Edinburgh University Press, 1971.
126. PITTS, W. H. AND W. S. McCULLOCH. How we know universals, the perception of auditory and visual forms. Bull. Math. Biophys. 9: 127-147, 1947.
127. PURDY, W. C. The Hypothesis of Psychophysical Correspondence in Space Perception (Ph.D. Thesis). Cornell University. Ann Arbor: University Microfilms No. 58-5594, 1958.
128. RAIBERT, M. H. A model for sensorimotor learning and control. Biol. Cybernetics 29: 29-36, 1978.

129. RANCK, J., Jr. Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. I. Behavioral correlates and firing repertoires. Exp. Neurol. 41: 461-522, 1973.
130. ROSENBLUETH, A., N. WIENER AND J. BIGELOW. Behavior, purpose and teleology. Philos. Sci. 10: 18-24, 1943.
131. ROSENFELD, A., R. A. HUMMEL AND S. W. ZUCKER. Scene labelling by relaxation operations. IEEE Trans. Systems, Man and Cybernetics SMC-6: 420-433, 1976.
132. SACERDOTI, E. D. Planning in a hierarchy of abstraction spaces. Artif. Intell. 5: 115-135, 1974.
133. SCHEIBEL, A. B. AND M. E. SCHEIBEL. Structural substrates for integrative patterns in the brain stem reticular core. In: Reticular Formation of the Brain, edited by H. Jasper et al. Boston: Little, Brown, 1958, p. 31-55.
134. SCHMIDT, R. A. A schema theory of discrete motor skill learning. Psychol. Rev. 82: 225-260, 1975.
135. SCHMIDT, R. A. The schema as a solution to some persistent problems in motor learning theory. In: Motor Control: Issues and Trends, edited by G. E. Stelmach. New York: Academic Press, 1976, p. 41-65.
136. SHIK, M. L. AND G. N. ORLOVSKII. Neurophysiology of locomotor automatism. Physiol. Rev. 56: 465-501, 1976.
137. STELMACH, G. E., Editor. Motor Control: Issues and Trends. Academic Press, 1976.

138. SZENTÁGOTHAJ, J. The modular organization of the cerebral cortex. Proc. Roy. Soc., 1978.
139. SZENTÁGOTHAJ, J. AND M. A. ARBIB. Conceptual Models of Neural Organization. Neurosci. Res. Program Bull. 12: 307-510, 1974.
(Republished by MIT Press, 1975.)
140. TOLMAN, E. C. Cognitive maps in rats and men. Psychol. Rev. 55: 189-208, 1948.
141. TREVARTHEN, C. B. Two mechanisms of vision in primates. Psychol. Forsch. 31: 299-337, 1968.
142. TURING, A. M. On computable numbers with an application to the Entscheidungsproblem. Proc. London Math. Soc. Ser. 2 42: 230-265, 1936.
143. WALTZ, D. L. A parallel model for low-level vision. In: Computer Vision Systems, edited by A. R. Hanson and E. M. Riseman. New York: Academic Press, 1978, p. 175-186.
144. WIENER, N. Cybernetics. Wiley/The Technology Press, 1948. (2nd edition, The MIT Press, 1961.)