TRANSFORMATIONS AND SOMATOTOPY
IN PERCEIVING SYSTEMS

Michael A. Arbib
Technical Report 70C-5
Department of Computer and Information Science
University of Massachusetts
Amherst, Massachusetts 01002

A paper for presentation at the Second International Joint Conference on Artificial Intelligence, Imperial College, London, England, September 1-3, 1971.

TRANSFORMATIONS AND SOMATOTOPY IN PERCEIVING SYSTEMS*

Michael A. Arbib
Department of Computer and Information Science
University of Massachusetts
Amherst, Massachusetts 01002 U.S.A.

Descriptors: Brain theory, robotics, ballistic and tracking movements, transformations, somatotopy, pattern recognition.

ABSTRACT: We present the notion of distributed computation in a layered somatotopically organized computer, present the Pitts-McCulloch scheme for obtaining standard forms, provide anarchic networks for ballistic and tracking modes of behaviour, and relate this to the visuomotor activity of the frog.

1. INTRODUCTION TO SOMATOTOPY

Much research in artificial intelligence seeks efficient ways to implement certain "intelligent" activities on computers, with little concern for the correspondence between the resultant mechanisms and those of the human brain. The present paper, on the contrary, belongs to that line of research which designs its artefacts as models to be used in increasing our understanding of brain mechanisms. By this we preclude attempts to implement pattern-recognition, say, in the manner which is most efficient on existing machines, but may nonetheless hope that our studies will offer clues for the design of future, highly parallel, computers for use in the control systems of robots.

To give a concrete example of our emphasis on parallelism, consider the problem of modelling the visual tracking of a moving object. Many authors, in applying control theory to the special case, say, of fixing the gaze upon a stationary or slowly moving object would note that two crucial parameters were involved - the present angle of gaze θ, and the desired angle of gaze, θ_d. They would then analyze the problem in terms of such a control system as shown in Figure 1, asking what function of the desired and actual gaze is computed to determine the rotational acceleration $\ddot{\theta}$ of the eye.

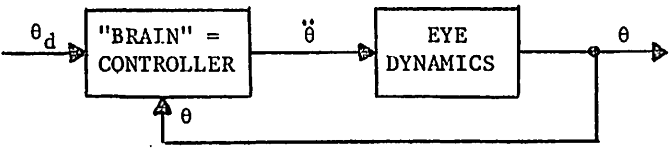

Figure 1

Such an approach has proved fruitful in analyzing behavior of biological systems, but may be dangerously misleading when it comes to unravelling the details of neural circuitry, for it suggests that we view the brain in terms of a central executive which manipulates a few variables such as θ and θ_d to issue such directives as the current value of $\ddot{\theta}$.

However, θ_d is not immediately available to the brain, but is instead encoded in terms of peaks of activity in a whole layer of neurons--the rods and cones of the eye. Again, in the case of eye dynamics, $\ddot{\theta}$ cannot be effective as a single control signal for a rotary actuator but must rather control the opposed activities of at least one agonist-antagonist pair of muscles - and even here, two signals are not enough, for the contraction of each muscle, itself a population of muscle fibres, must result from the overall activity of a whole population of motoneurons.

Thus, although it would be possible to design a robot with a "brain" structured like the centralized $(\theta_d,\theta) \mapsto \ddot{\theta}$ converter of Figure 1 it would require special preprocessors to "funnel down" the whole input array of retinal activity to provide the single number θ_d. Indeed, this scheme might make sense in a robot whose task was to track single targets rather than interact with complex environments, and whose effector was a single rotary actuator for which $\ddot{\theta}$ was an appropriate control signal. But if the output must be played out upon a whole array of motoneurons, as in the biological case, so that $\ddot{\theta}$ would have to be fed into an elaborate processor to be "parcelled out", then one begins to doubt the utility of the centralized processor. In Section 2 we shall recall a scheme, due to Pitts and McCulloch, whose beauty lies in the simplicity of its demonstration that - at least in the case under discussion - a centralized processor may be dispensed with, and all computation may be carried out in distributed fashion in the layer or layers between the input and output arrays. In Section 4, we shall outline a related model of frog visumotor activity. Some data about the frog visual system may help make our point:

Lettvin, Maturana, McCulloch and Pitts (1) found that most ganglion cells of the frog's retina could be classified as being one of four types - such as "moving spot detectors" and "large moving object detectors". What we want to emphasize here is the way in which the information from the four types of detectors is

---

*Preparation of this report was supported in part by the U.S. Army Research Office Durham under contract No. DAHC 04-70-C-0043 and in part by the Public Health Service under grant No. NS 09755-01 COM from the National Institute of Neurological Diseases and Stroke, and in part by the University of Massachusetts under Faculty Research Grant S12-71(1).

distributed in the brain. Their axons terminate (among other places) in a brain structure called the <u>tectum</u>, with the terminations forming four separate layers, one atop each other, with the properties that (a) different layers correspond to different types of detector; (b) each layer preserves the spatial relations between the original cells (i.e. there exists a direction along the layer corresponding to moving across the retina); (c) terminations stacked above one another in the four layers come from ganglion cells with overlapping receptive fields. This is another dramatic case of the neural specificity that provides the structural substrate for brain function (see also Sperry (2)). It should be noted that such relationships between two layers may preserve rough spatial relationships (up and down versus across), without preserving relative sizes. For example, in the layer in the human brain which receives touch information from the body, the fingers occupy a larger area than the trunk, since the brain needs detailed sensory information from the fingers if it is to control fine manipulation. Such relationships between two layers of cells are called <u>somatotopic</u>, from the Greek <u>soma</u> (body) and <u>topos</u> (place), since it preserves information about place on the body as we move from receptors to the central nervous system. As we move further from the periphery, the relationships become less distinctive, but may still guide our investigation of adjacent layers. What we are saying is that a useful way to structure the apparent chaos of many parts of the brain is to describe such parts in terms of interconnected layers, where position within the layer is a crucial indicator of the functional significance of a cell's activity, and where an analysis of one patch of such a layer may yield an understanding of the function of the layer as a whole.

In discussing somatotopy in the layers of such a distributed computer, the reader should take note that we shall use the word somatotopy in an extremely broad fashion. In the input pathways of the visual system, position encodes position in visual space relative to the eye; in the auditory system it encodes frequency of stimulation; and in the tactile system, position on the body. It is only in the last case that the term somatotopy is strictly appropriate - retinotopy and tonotopy may better connote the respective situations in the first two cases. Again, in output pathways, position in a layer may encode the location of the target of a movement. As we move away from the periphery to layers of the brain far removed from any predominant commitment to sensory modality on particular mode of action, we can expect position in the layer to have little direct correlation with bodily position - yet we hypothesize that position in the layer will still encode a crucial parameter of the cell's function. It is in this somewhat over-extended sense of a positional code that we shall speak of somatotopy even in layers far from the periphery. There are structures in the brain - the reticular formation may be one - where the notion of layering is not useful.

We do not fanatically claim the universal truth of the statement "the brain is a layered somatotopic computer". Rather we use it as a convenient slogan to remind us that it is high time that somatotopy - so long an important property for anatomists and physiologists - played its full role in our theories of brain function. Even in structures which are not layered, the positions of neurons will play a role that we cannot neglect in modelling their contribution to the overall function of the structure.

2. PITTS AND McCULLOCH REVISITED

In the latter part of their paper, Pitts and McCulloch (3) presented a feedback scheme designed to find a transformation T, from among a group G of possible transformations of patterns which are played, say, upon the retina, which will transform a pattern ϕ to a standard form $\phi_0 = T\phi$ [Arbib and Didday (4) consider the case in which we also use transformations of output activity, to assure that the relation between input and output is in standard form.]

They generate the transformation in two steps:

(i) Associate with each pattern ϕ an "error vector" $E(\phi)$ such that $E(\phi) = 0$ if and only if ϕ is in standard form.

(ii) Provide a scheme $\mathcal{W}$ which will associate with each error vector a transformation which is error-reducing - that is, for all patterns ϕ we demand that $E(\phi)$ be reduced after $\mathcal{W}(E(\phi))$ is applied to ϕ:

$$|| E [\mathcal{W}(E(\phi))\cdot\phi] || \leq E(\phi) \qquad (1)$$

with equality only in case $E(\phi) = 0$. Henceforth, let us use W_ϕ to abbreviate $\mathcal{W}(E(\phi))$.

There are two main implementations of such a feedback scheme, only the second of which was considered by Pitts and McCulloch.

In a <u>ballistic scheme</u>, $\mathcal{W}$ is so structured as to virtually reduce the error to zero in one step:

$E[W_\phi \cdot \phi] \doteqdot 0$ for all patterns ϕ

A controller would then proceed as follows:

1. Given ϕ, compute $E(\phi)$ and thus W_ϕ.
2. Form $W_\phi \cdot \phi = \hat{\phi}$
3. Proceed on the assumption that $\hat{\phi}$ is sufficiently close to standard form.

Such a scheme is that used in ballistics where $E(\phi)$ is the displacement of a bullet from its target, and W_ϕ is determined by the initial aim when the shot is fired - there is no possibility of making mid-course corrections. This is in distinction to a guided missile in which repeated corrections can be made.

In a tracking scheme, then, the error $E[W_\phi \cdot \phi]$ may be little less than the previous error $E[\phi]$ - all we demand is that under a feedback scheme employing repeated application of the error-correction the error eventually go to zero. A controller implementing tracking may proceed according to one of two schemes. The implementation mode corresponds to continually modifying the pattern until one is found which is in standard form; the planning mode corresponds to continually modifying the transform until one is found which will bring the given pattern to standard form:

```
I. Here $\phi$ will be the latest transformed version
   of the input pattern.
   1. Replace $\phi$ by the new input pattern.
   2. Use $E(\phi)$ to obtain $W_\phi$
   3. Form $W_\phi \cdot \phi$ to obtain the new pattern $\phi$
   4. Is the new $E(\phi)$ close enough to zero?
      YES: Exit, $\phi$ may be treated as in standard form
      NO:  Go to 2
II. Here $\phi$ will be the fixed input pattern, and
    $T$ will be the updated transform to be applied to $\phi$
   1. Initialize $T$ to be the identity transformation: $I\phi = \phi$
   2. Use $E(T\phi)$ to obtain $\mathcal{W}[E(T\phi)]$
   3. Form $\mathcal{W}[E(T\phi)] \cdot T$ to obtain the new transform $T$
   4. Is the new $E(T\phi)$ close enough to zero?
      YES: Exit, $T\phi$ may be treated as in standard form
      NO:  Go to 2
```

To guarantee that the schemes converge we need a stronger condition than (1) - one condition is that there exists some number δ such that $0 < \delta < 1$ and

$$|| E(W_\phi \cdot \phi) || \leq (1-\delta) || E(\phi) || \text{ for all patterns } \phi.$$

Convergence then follows from the fact that $(1-\delta)^n \to 0$ as $n \to \infty$.

There will be applications in which a controller may wish to use a mixed ballistic-tracking strategy - using a transform generator $\mathcal{W}_1$ to compute a first "giant leap" to bring the pattern fairly close to standard form, then a second transform generator $\mathcal{W}_2$ to be used in a tracking strategy to iteratively "fine tune" the pattern ever closer to standard form. In fact, this "combined strategy" seems to be that employed in many biological systems (5).

Figure 2 shows a discrete-time system (6) which will implement Scheme II to generate, for any ϕ, a transformation T_ϕ which will transform it to standard form.

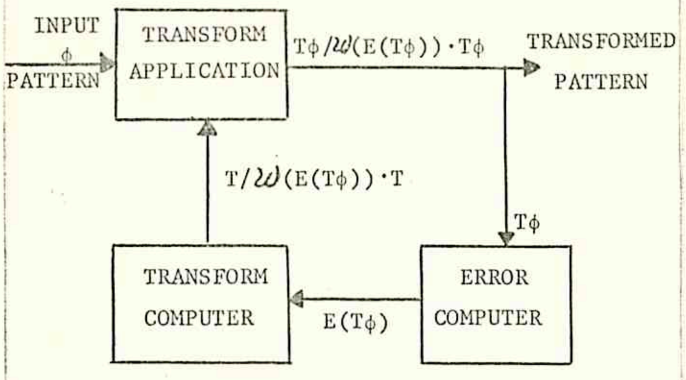

Figure 2

A Generalization of the Pitts-McCulloch Scheme for Transforming a Pattern to Standard Form. (6)

The transform application box is memoryless-input pattern ϕ and transform T at its input yield transformed pattern $T\phi$ at its output. The error computer box is memoryless - an input pattern at its input yields the corresponding error at its output. The transform computer box is a sequential machine - if its state at time t is the transform T, and its input at time t is the error vector e, then its new state and output at time $t + 1$ will both be the transform $\mathcal{W}(e) \cdot T$.

The hard work in such a scheme is actually defining an appropriate error measure E and then finding a mapping which can make use of error feedback to properly control the system so that it will eventually transform the input to standard form.

Let us see here how Pitts and McCulloch exemplified their general scheme in a plausible reflex arc from the eyes through the superior colliculus to the oculomotor nuclei to so control the muscles which direct the gaze as to bring the point of fixation to the centre 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 above prescription. But that does not affect the model which follows].

Julia Apter (7,8) showed that each half of the visual field of the cat (seen through the nasal half of one eye and the temporal half of the other) maps topographically upon the contralateral colliculus. In addition to this "sensory" map, she studied the "motor" map by strychinizing 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.

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 centred when excitation from the left is exactly balanced by excitation from the right. Their model (3, Figure 6) is then so arranged, for example, that each motoneuron controlling muscle fibres in the left medial rectus and right lateral rectus muscles, which contract to move the left and right eyeballs, respectively, 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 the right lateral and medial rectus, i.e., the 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-centre of gravity's vertical coordinate is done similarly. [Of course, computation may be performed by commisural fibres linking similar contralateral tectal points, instead of in the oculomotor nuclei.] Eye movement ceases when and only when the fixation point is the centre of gravity.

It must be emphasized that the reflex for which we have just summarized a crude, though instructive, model would be subject to "higher control" in normal function (5). For example, "interest" might be the criterion for determining which area of the visual field to examine, with the reflex determining the fixation point within the region (cf. the fine tuning servo on a radio receiver) - gaze may then remain fixed at that point until it is "adequately" perceived. Conversely, a sudden flash may usurp the averaging operations to dominate the reflex control of gaze momentarily, forcing the organism to attend at least briefly to a novel stimulus.

In the context of this conference, it seems well worth noting that the general scheme of Figure 2 is essentially that suggested independently more than a decade after (3) by Newell, Shaw and Simon (9) for their GPS (General Problem Solver). The correspondence between GPS' general framework for solving problems and the Pitts-McCulloch scheme is as follows:

1. GPS is given a set of objects [This corresponds to our set of patterns].
2. GPS is given a finite set of differences, and a means to determine which differences obtain between a pair of objects. [This corresponds to our error function E, but since the set of differences is finite, can only give rough indications as to "what is wrong"].
3. GPS is given a finite set of operators, and an operator-difference table, which lists for each difference the operators likely to reduce it. [This corresponds to our transform generator $\mathcal{W}$].
4. GPS is given an initial object (say the list of axioms in a propositional logic) and a final object (say a statement we should like to prove to be a theorem) and is to find a sequence of operators which will transform the initial object into the final object. (If the operators correspond to rules of inference, then in our example, the desired chain of operators would provide the desired proof of the given statement.)

The catch here is that since the difference only gives us very partial information about what needs to be changed, we cannot guarantee that applying a recommended operator will indeed transform the latest object into one that is genuinely closer to our goal. Further, a given operator will not be applicable to all objects, and so preliminary transformations may be required to place an object in a form to which an indicated operator can be applied. Because of this one cannot proceed a step at a time as in the Pitts-McCulloch scheme. Rather, one must develop a "decision tree" in which we keep track of the application of various possible operators at various stages. The aim of the general supervisory part of the GPS program is to ensure that we put most effort into "growing" those branches of the tree which seem to be leading towards the goal. For each node we could find the differences between it and the goal, and then determine which operators are suggested by the difference table. Some of them may not be applicable. At each stage we must decide which is the most promising node to next operate upon, and which of the possible operators we should apply.

Thus GPS involves a supervisory control with memory of various paths which may yet be found to lead to the goal - it is the design of this control program that really sets GPS off from the Pitts-McCulloch scheme. [An approach which merges such considerations of "heuristic search" with an application closer to that of (3) may be found in (4).] The claim of GPS to generality is that it can solve any problem - such as proving theorems in propositional logic - which can be solved using a tree-search on the basis of an operator-difference table. Unfortunately, this very generality makes for the inefficiency which attends the inability to use special tricks developed for a given problem domain. Further, it must be stressed that not all problems are amenable to this type of solution. Even for problems which are, the real intelligence usually comes not so much in using a given operator-difference table, but rather in realizing what differences are salient features of the given problem and generating by experience a table of the operators likely to reduce them. Perhaps techniques akin to those Uhr and Vossler have used in programming a machine to generate its own set of feature detectors for pattern recognition may eventually be developed to provide an operator-difference table generating routine to complement the supervisory package of GPS.

3. DISTRIBUTED MOTOR CONTROL

It should be noted that even if the mathematical equations formalizing the Pitts-McCulloch scheme of Figure 3 for centering of gaze were to contain a damping term to prevent the eyeball from undergoing continual oscillations, it still has the defect of being essentially a tracking model, whereas the reflex "snapping" of gaze toward a flash of light is essentially ballistic. In fact, human eye movements can be either ballistic or tracking. Typically, a human examining a scene will fixate on one point of the visual field then make a saccadic movement (the term for a ballistic eye movement) to fixate another point of the visual scene, until satisfied that he has scanned enough of the scene to perceive his current environment. However in other situations - such as watching a car go by before crossing the street - he will fixate upon an object and then track it. In man, various cortical areas can modulate activity in superior colliculus, and Bizzi has found in one of them - the so-called frontal eye field, which is in frontal cortex - that there are three types of cells, type I which are active in ballistic eye movements, type II which are active in tracking eye movements, and other cells more concerned with head movements than with eye movements. Perhaps a similar situation will be found on closer examination of superior colliculus. In any case, it does seem that the Pitts-McCulloch scheme is more suited to the tracking mode than to the ballistic mode. To rectify this, let us then present another model, due to Braitenberg and Onesto (10) for a distributed computer controlling ballistic movement. (It should be mentioned that they conceived their model as a model of the cerebellum, but subsequent investigations have revealed so much new data about the cerebellum that their model cannot stand as a model of the cerebellar cortex without drastic modification. The reader may find a thorough critique of cerebellar modelling in Boylls and Arbib (11), but it would not seem fruitful to present the details here, for our aim in this paper is not to say "Here is the correct model for the function of a certain subsystem of the brain", but rather to say "Here is a fruitful way to go about modelling brain function". In this spirit, we present models which give one new principles of organization, hoping in this way to spur much further work to find the biological implementation of these principles in neural circuitry; or to see their refinement in the design of control circuitry for robots.)

When a shot is fired from a gun two forces are involved - the explosion that propels the projectile towards the target, and the braking force that results when the projectile hits the target (if the target were to step aside, the projectile would not stop in the position at which it was originally aimed). Ballistic movements in animals also involve this "bang-bang" control. There is an initial burst of acceleration as the agonist contracts and the antagonist muscle relaxes; an intervening quiet period; and then the final deceleration as the antagonist contracts. Experiments on rapid flexion and extension of joints have shown that muscle activation occupies only a small portion of the movement, that the duration of this activation does not seem to be related to the extent of the movement. Thus the duration of the movement seems to be determined mainly by the timing, relative to the "go" signal, of the "stop" signal (which has to be determined by the brain, rather than being imposed by the environment, as it was in our projectile example). Braitenberg and Onesto thus proposed a network for converting space into time (a subtle alchemy!) by providing that the position of an input (encoding the desired target position) would determine the time of the output (which would trigger the "slamming on of the brakes"). The scheme (see Figure 3) has a linear array of output cells whose output circuitry is so arranged that the firing of any one of them will yield the antagonist burst that will brake the ballistic movement. There are two systems of input fibres each arranged in the same linear order, with position along the line corresponding to angle of flexion of the joint. The first class, which we shall call the C-fibres, connect to a single output cell. The second class, which we shall call the M-fibres, bifurcate into fibres which contact each cell in the array. The speed of propagation along these parallel fibres is such that the time required to go from one point in the array to another corresponds to the time the joint requires to move between the corresponding angles.

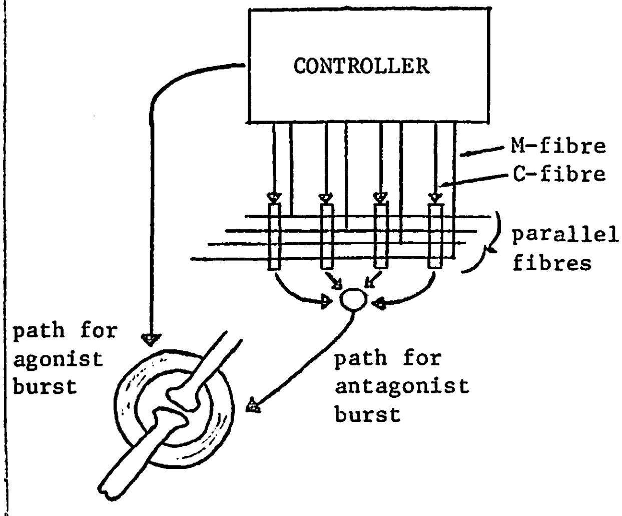

Figure 3

The Braitenberg-Onesto Scheme for Control of Ballistic Movement

The controller then elicits a ballistic movement by firing 3 signals - one to trigger the agonist burst which will initiate movement, one on the C-fibre corresponding to the initial joint position, and one on the M-fibre corresponding to the target position. If we assume that an output cell can only respond to parallel fibre input if it has received C-fibre input, we see that only the output cell corresponding to the activated C-fibre will fire, and it is clear that its time of firing will correspond to its distance from the activated M-fibre. Thus it will elicit the braking effect of the antagonist burst at precisely the right time.

Note that in the above scheme, we could relieve the controller of having to "know" where the joint is, by having a feedback circuit continually monitor joint position and keep the appropriate C-fibre activated.

While we do not claim to have modelled the way the nervous system controls movement, what we have shown is that a plausible subsystem for vertebrate nervous systems may be of the type shown in Figure 4 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 "centre of gravity" of the positions encoded by that array.

It should be noted that a full elaboration of this scheme would involve hierarchical arrangements. For example, in fixating a new point in space, 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, whose behavior would of course be modified by the low-level postural controllers in the brainstem and spinal cord. We should also add that the scheme must be elaborated to provide for generating particular velocities, etc. To caricature it crudely, one may conjecture that such an option has evolved through the development of circuitry which can control tracking movements internally, rather than driving them through sensory channels.

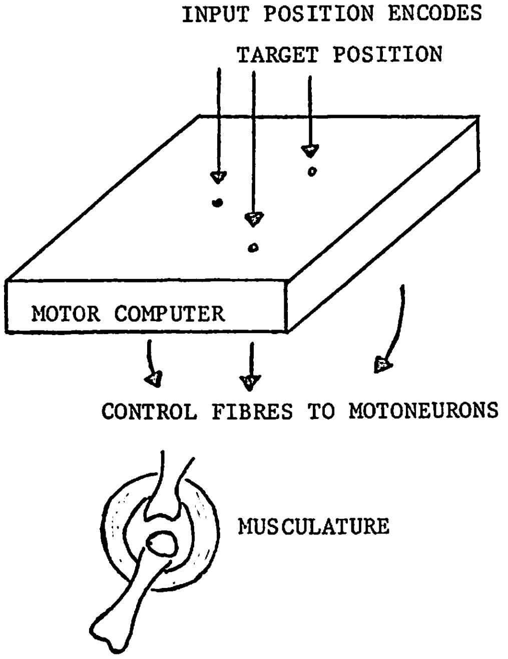

Figure 4

Schematic for Distributed Motor Control

4. TWO EXAMPLES

To round out our discussion, let us present a model of frog visumotor behaviour which involves layered distributed computation, and then close by relating it to a model of the reticular formation. First, we need to comment on the idea of "Redundancy of Potential Command". If we take the position that perception of an object generally involves the gaining of access to "programs" for controlling interaction with the object, rather than simply generating a "label" for the object, we must emphasise gaining of access to a program rather than the execution of a program - one may perceive something and yet still leave it alone. Thus in gaining access to the program, the system only gives it potential command, further processing being required to determine whether or not to act. A key question will thus be "How is the central nervous system structured to allow coordinated action of the whole animal when different regions receive contradictory local information?" McCulloch suggested that the answer lay in the Principle of Redundancy of Potential Command which states, essentially, that command should pass to the region with the most important information. He cited the example of the behavior of a World War I naval fleet controlled - at least temporarily - by the signals from whichever ship first sighted the enemy, the point being that this ship need not be the flagship, in which command normally resided.

McCulloch further suggested that this redundancy of potential command in vertebrates would find its clearest expression in the reticular formation of the brain stem (RF). Kilmer and McCulloch then made the following observations towards building a model of RF:

(i) They noted that at any one time an animal is in only one of some 20 or so gross modes of behavior - sleeping, eating, grooming, mating, urinating, for example - and posited that the main role of the core of the RF (or at least the role they sought to model) was to commit the organism to one of these modes.

(ii) They noted that anatomical data of the Scheibels (12) suggested that RF need not be modelled neuron by neuron, but could instead be considered as a stack of "poker chips," each containing tens of thousands of neurons, and each with its own nexus of sensory information.

(iii) They posited that each module ("poker chip") could decide which mode was most appropriate to its own nexus of information, and then asked, "How can the modules be coupled so that, in real-time, a consensus can be reached as to the mode appropriate to the overall sensory input, despite conflicting mode indications from local inputs to different modules?"

This was the framework within which Kilmer, McCulloch and Blum (13) designed and simulated the compartment model, called S-RETIC, which we have discussed above of a system to compute mode changes, comprising a column of modules which differed only in their input array, and which were interconnected in a way suggested by RF anatomy.

Pitts and McCulloch's model (our Section 2) of the superior colliculus (which is the cat's "equivalent" of the frog's tectum) was offered as a plausible explanation of how an animal might fixate its gaze at the "average" or "centre of gravity" of a field of illumination. For us, their scheme has the added significance that it showed how to design a somatotopically organized network in which there is no "executive neuron" which 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.

If we paraphrase our interpretation of the significance of the Pitts and McCulloch model of the superior colliculus to say that it showed how "the organism can be committed to an overall action by a population of motoneurons none of which had global information as to which action is appropriate", we are struck by the similarity of the situation to that in our statement of the RF problem.

We may build on this to illuminate another system for the study of redundancy of potential command. The frog, which is normally immobile, will snap at any fly that comes into suitable range - "snapping" comprising a movement of the head (and, when necessary, the body) to aim at the fly and the rapid extension of the tongue to "zap" the fly. The situation seems very simple in that the frog does not seem to recognize flies as such - rather it will snap at any wiggling object, but will not snap at a stationary (i.e. dead) fly. A frog confronted with two flies then presents us with a beautifully simple redundant command situation - normally the animal snaps at one of the flies, and so we have sought to model the brain mechanism that determines which of the flies will "take command" of the frog. This could be explained in terms of a serial scan made of the tectum until a region is first found in which the activity in the four layers of ganglion cell termination in the tectum signals the presence of a fly - at which stage the scanner would issue a command to snap in the direction indicated by the current address of the scan. However, we argued that such serial processing is not a candidate for the frog's neural machinery because of the fact (among others) that the frog will sometimes snap midway between two flies - precisely the "center of gravity" effect one expects from an output system of the distributed computation type suggested by Pitts and McCulloch for centering of gaze. [Note that the above distinction between serial and distributed processing could not be made by asking only the usual question of sensory physiology, "What information is relayed to the brain?" but by also asking, "How does the animal make use of such information to act?"]

However, we must note that while the Pitts-McCulloch model does yield integrated behavior, it does not explain the "usually-one-fly-effect." Didday (14,15) has offered a mechanism for this which, in retrospect, could be seen to bear a great resemblance to the Kilmer-McCulloch RF model. The observations on frog behavior suggest three layers of processing, each involving distributed computation. The first layer operates upon the four layers of retinal information to provide for each region a measure of "foodness." The third layer does a modified Pitts-McCulloch type computation to direct motion of the frog to the position corresponding to the "center of gravity" of activity in the second layer. The task of the second layer is then very much akin to the task of the Kilmer-McCulloch RF. Where that model (13) has an array of modules which must interact to get a majority favoring the same mode, the task of the second layer of our hypothetical tectum (15) is to turn down the activity of all but one region of (or from) the first layer. The essential mechanisms turn out to be very similar, and provide a plausible analogue for the "sameness" and "newness" neurons observed by Lettvin et al. (1). The models differ in having all modes evaluated in each module, versus having a module identified with a mode. In any case, the study of frog behavior sheds new insight on RF modelling, and suggests alternate hypotheses. Our model is still a crude oversimplification of the complexities of a real frog brain, but we believe that our partial successes show that our organizational principles, all too often neglected in the cybernetics literature, must play a crucial role in future brain theory. This thesis is elaborated at book length in (5), and extended in (4) and (11). See also the related studies of Greene (16-19).

REFERENCES

(1) J. Y. Lettvin, H. R. Maturana, W. S. McCulloch and W. H. Pitts (1959): "What the Frog's Eye Tells the Frog's Brain," *Proc. IRE*, 47, 1950-1959.

(2) R. W. Sperry (1951) "Mechanisms of Neural Maturation," in *Handbook of Experimental Psychology*, New York: John Wiley, 236-280.

(3) W. H. Pitts and W. S. McCulloch (1947): "How We Know Universals: The Perception of Auditory and Visual Forms," *Bull. Math. Biophys.*, 9, 127-147.

(4) M. A. Arbib and R. L. Didday (1971): "The Organization of Action-Oriented Memory for a Perceiving System. Part I. The Basic Model," *Journal of Cybernetics* 1, No. 1. (Part II, "The Hologram Metaphor," by M. A. Arbib and P. Dev, will appear in a later issue.)

(5) M. A. Arbib (1972): *The Metaphorical Brain*, New York: Wiley-Interscience.

(6) M. A. Arbib (1971) "How We Know Universals: Retrospect and Prospect," *Mathematical Biosciences*, Warren McCulloch Memorial Issue.

(7) J. Apter (1945): "The Projection of the Retina on the Superior Colliculus of Cats," *J. Neurophysiol.*, 8, 123-134.

(8) J. Apter (1946): "Eye Movements Following Strychinization of the Superior Colliculus of Cats," *J. Neurophysiol.*, 9, 73-85.

(9) A. Newell, J. C. Shaw and H. A. Simon (1959) "Report on a General Problem-Solving Program" Proc. International Conf. Info. Proc., Paris: UNESCO, 256-264.

(10) V. Braitenberg and N. Onesto (1960): "The Cerebellar Cortex as a Timing Organ," *Congress Inst. Medicina Cibernetica*, First Naples Atti., 239-255.

(11) C. C. Boylls and M. A. Arbib: "Organizational Principles in Cerebellar Modelling," *Progress in Biophysics*, (to appear).

(12) A. B. Scheibel and M. E. Scheibel (1958): "Structural Substrates for Integrative Patterns in the Brain Stem Reticular Core," in *Reticular Formation of the Brain* (H. H. Jasper et al., Eds.), Boston: Little, Brown, 31-55.

(13) W. L. Kilmer, W. S. McCulloch and J. Blum (1968): "Some Mechanisms for a Theory of the Reticular Formation," in *Systems Theory and Biology* (M. Mesarovic, Ed.), New York: Springer-Verlag, 286-375.

(14) R. L. Didday (1971 a) "An Easy Method for Simulating Distributed Computation in Nervous Systems," *Int'l J. of Man Machine Studies*. (in press)

(15) R. L. Didday (1971 b) "A Possible Decision-Making Role for the 'Sameness' and 'Newness' Cells in the Frog," submitted to *Brain Behav. and Evol.*

(16) P. H. Greene (1964) "New Problems in Adaptive Control," in *Computer and Information Sciences* (J.T. Tou and R.H. Wilcox, Eds.), Washington: Spartan.

(17) P. H. Greene (1967) "Models for Perception and Action", *Proceedings of the First Annual Princeton Conference on Information Sciences & Systems*, Dept. of EE, Princeton University, 245-253. In relation to which, see numerous articles by Russian workers in *Biophysics*, the English translation of the Russian *Biofizika*, in the Section entitled "Biophysics of Complex Systems. Mathematical Models."

(18) P. H. Greene (1968) "A. Essential Features of Purposive Movements; B. Cybernetic Problems of Sensorimotor Structure: Introductory Remarks and Survey of a Study; C. Seeking Mathematical Models of Skilled Actions; D. An Aspect of Robot Control and Nervous Control of Skilled Movements; Coordination of Two Effectors and Transfer of Adaptation," Institute for Computer Research, Quarterly Progress Report No. 16, Sections III A-D.

(19) P.H. Greene (1970) "The Theory of Tasks: Cybernetic Problems of Sensorimotor Structure," book in preparation (Preliminary material may be found in Greene [1967, 1968]).