Organizational Principles for Theoretical Neurophysiology[1]

Michael A. Arbib
Department of Computer and Information Science
Technical Report 70C-6

University of Massachusetts
Amherst, Mass. 01002

(August 15, 1971)

A contribution to "Towards a Theoretical Biology 4",
C.H. Waddington, ed., Edinburgh: Edinburgh University Press (1972).

---

[1] This paper has grown out of brief remarks made at Serbelloni IV, and is based on two other conference papers: "Organizational Principles for Embryological and Neurophysiological Processes" presented at a conference on Physical Principles of Neural and Organismic Behaviour [held 16-18 December 1970 at the Center for Theoretical Studies, University of Miami, Coral Gables, Florida]; and "Transformations and Somatotopy in Perceiving Systems" presented at the Second International Joint Conference on Artificial Intelligence [held 1-3 September 1971 at Imperial College, London]. The arguments of this paper are elaborated in the author's book "The Metaphorical Brain" (Wiley-Interscience, 1972). Preparation of this paper was supported in part by the Public Health Service under Grant No. 2-R01 NS 09755-02 COM from the National Institute of Neurological Diseases and Stroke, and in part by the U.S. Army Research Office Durham under Contract No. DAHC 04-07-C-0043.

CONTENTS:

1. *The "Top-Down" Approach*: We suggest the need for organizational principles to complement the study of molecular mechanisms in theoretical biology.

2. *Introduction to Somatotopy*: We consider a simple model of control of eye movements to suggest somatotopy as an ordering principle for the description of neural architecture.

3. *Neural Specificity and Theoretical Embryology*: We stress that the neonate's brain is no *tabula rasa*, and briefly discuss the problems of theoretical embryology raised by the requisite neural specificity.

4. *A Distributed-Computation Model for the Control of Gaze*: We present a class of feedback schemes, each designed to find a transformation which will bring a pattern to standard form, and give as example Pitts and McCulloch's model of the superior colliculus as a controller of gaze.

5. *Distributed Motor-Control*: The Pitts-McCulloch model is contrasted with the Braitenberg-Onesto model for the control of ballistic movements to suggest a general strategy for distributed control of movement.

6. *Redundancy of Potential Command*: We contrast the decision-making of a frog in snapping at one of several flies with the process whereby a vertebrate is hypothesised to choose one of several gross modes of action.

7. *Organizational Principles*: We summarize and extend the above discussion by presenting eight organizational principles for theoretical neurophysiology.

1. THE "TOP-DOWN" APPROACH

By now, many hitherto mysterious properties of cells - the basic "components" of all organisms - have been explained in terms of biochemistry and molecular biology, and many papers document the power of such explanations. It would be foolish to try to belittle such achievements, but it would be equally mistaken to be so dazzled by their success as to believe that biochemistry alone can unravel all the knotty problems of biology, and that the development of new theoretical approaches is unnecessary.

Going "up" from the basic biochemistry and physics must be complemented by going "down" from overall functional questions, as in the division of labour in Computer and Information Science between the component expert (using solid-state physics in trying to push to the ultimate reduction of size, increase in speed of operation and flexibility of function for devices which are then to be built into computers) and the computer scientist (who studies how to put together large scale organizations in terms of components of stipulated function to get some overall sophisticated behaviour). Similarly, in biology the component, or cellular, level is the meeting ground for two quite different approaches. To explain how cells "work", and their capacities for interaction, is the task of the biophysicist and biochemist; while explaining how to organize large collections of such components seems to require such approaches of Computer and Information Science as automata theory and computer simulation.

To claim this is not to claim that the right tools for a theoretical biology are readily at hand. One of the greatest pitfalls facing the mathematician, engineer, computer scientist or physicist turning theoretical biologist is that too much of his education has involved his mastering long-polished mathematical techniques, and finding that wide classes of problems can be solved simply by "plugging-in" these techniques. With this as background, it is all too easy to believe that he can solve the biologist's problems by "plugging-in" these techniques to biological situations, little realizing that a great period of induction and experimentation (for even theorists must experiment, even if only with symbolic constructs) was required to match technique to problem. But theory _is_ required in biology - as in any science where constructs become subtle enough to escape the domain of the immediately observable and where the depth of argument comes to exceed the usual grasp of common sense - and what remains to be determined is not whether there shall be theoretical biology, but rather what forms theoretical biology shall take. The theorist who can make a substantial contribution will probably be one who combines an intimate knowledge of the experimental data of some restricted problem in biology with a broad command of theoretical techniques, and uses the interaction between his reformulation and reconceptualization of the data and his reworking of the techniques to evolve genuinely new insights into that particular biological problem - only to find that those insights are valid elsewhere. There is no recipe for this, but in this paper we shall at least present some organizational principles that anatomical, physiological, behavioral and simulation studies suggest may provide a useful framework for theoretical neurophysiology.

2. INTRODUCTION TO SOMATOTOPY

To describe a neural network, we must discern order in it complexity. Here we suggest somatotopy as an ordering principle of neural architecture. To see the theoretical implications of this suggestion, we start with a simple consideration of the control of eye movements.

Consider the problem of applying control theory to the fixing of gaze upon a stationary or slowly moving object. Many authors 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 ask what function of the desired and actual gaze is computed to determine the rotational acceleration $\ddot{\theta}$ of the eye.

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 as a centralized $(\theta_d,\theta) \mapsto \ddot{\theta}$ converter 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 4 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 6, we shall outline a related model of frog visumotor activity.

Some data about the frog visual system may help make our point about the importance of somatotopy in neural design: Lettvin, Maturana, McCulloch and Pitts [1959] 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 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.

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 some parts of stimulation; and in the tactile system, position on the body. It is only in the last case that the term somatotopy (from the Greek soma (body) and topos (place)) is strictly appropriate, since it preserves information about place on the body as we move from receptors to the central nervous system - 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.

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.

As we move away from the periphery to layers of the brain far removed from any predominant commitment to sensory modality or 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. 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. 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.

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 reticular formation may be one - the positions of neurons will play a role that we cannot neglect in modelling their contribution to the overall function of the structure.

3. NEURAL SPECIFICITY AND THEORETICAL EMBRYOLOGY

As our discussion of somatotopy in Section 2 has emphasised, there is a great deal of specificity in neural architecture. In this section we shall present further examples of such specificity (for a full review, see Jacobson [1970]) and hint at some of the modelling problems which they raise for theoretical embryologists.

The stimulus-response view of behaviour held that if a creature, with adequate receptors and effectors, were put in some complicated environment and "punished" when it did something "wrong" and "rewarded" when it did something "right", then eventually the correct connections would be made to enable the organism to function effectively in that environment. Translating this into neural terms, many people thought the nervous system was completely "plastic": that all connections could be - and in fact were - moulded by experience. We owe to Paul Weiss, Roger Sperry and other workers the knowledge that plasticity is <u>not</u> unlimited, and that in fact there is a great deal of neuronal specificity - i.e. <u>genetics constrains many details of neuronal connections which cannot be changed by experience unless there exist specific brain structures to exploit that experience</u>.

A newborn baby has to be able to suck, to breathe, to excrete and so on, but it cannot do many other things at birth, and has to be able to learn how to do them. This cannot happen unless it has appropriate structures to implement learning. In other words, neural plasticity - paradoxically - requires a constraining neural specificity to be fully effective. This point may seem obvious, but is so often lost sight of that it may pay to belabour it: Think of tossing a coin repeatedly. Every time it comes up heads, spray it with Chanel No. 5, and every time it comes up tails spray it with stale cabbage juice. It hardly seems profound to doubt that the coin will eventually tend to come up heads rather than tails, but it may be helpful to explicate the grounds for our doubt.

Firstly, the coin does not have receptors which allow it to distinguish Chanel No. 5 from cabbage juice. Secondly, even if it could distinguish them, it has no inbuilt criteria to determine which is preferable. Thirdly, even if it could tell which was preferable it has no mechanism whereby it could make use of that knowledge to change its behaviour. Thus, in looking at the embryology of the nervous system we have to look for specificity, whether in the direct sense of determining networks which will mediate innate behaviour patterns, or to provide the adaptational substrate to enable the organism to adapt its evolutionary heritage to the exigencies of its own environment. We have to understand how appropriate receptor and effector arrays can be structured, how basic drive mechanisms can be "built into" the organism so that it can shape its behaviour on the basis of some evolutionarily determined criteria of biological usefulness or destructiveness, and we must understand - at least in mammals - the determination of a sufficiently rich cortical structure to allow sophisticated learning.

To enhance the latter point by a striking contrast, we may recall Paul Weiss' [1941, for an overview] intriguing experiments in which the forelimbs of a salamander were reversed in the larval stage. When such a salamander grew to salamanderhood, whenever it would see some food in front of it, the brain would send the appropriate command of "advance", but the neural circuitry in the brainstem which interpreted the command did not "know" that the forelegs were back-to-front and so would send the sequence of muscular activation which would cause the forelimbs to make the animal scurry away from its food. No matter how long the animal was exposed to this unfortunate situation it could never learn what was wrong - or, at least, if it learned what was wrong, it could not do anything about it. Thus we see the necessity for adequate structure if learning is to ensue.

Notice that what we are talking about in the nervous system is not the development of individual organs per se, but rather the development of functional systems which involve the whole organism. The animal at birth has to be able to take tactile stimuli on the lips and go through the "computation" required to convert this into a sucking reflex. If we look at animals such as the guinea pig in which the hindlimbs are more important than the forelimbs at birth we will find the uneven development of the spinal cord which insures that the hindlimbs are ready to function at birth. This is what the Russian physiologist Anokhin [1964] refers to as systemogenesis - we have to think of the nervous system not in terms of anatomically defined lumps of tissue, but rather in terms of an interacting overlapping collection of systems for carrying out biologically important functions. Thus, our task becomes even more complicated when we realize that it is not enough to look at one small part of the body or the nervous system and explain how it grows, but we have to explain the sort of synchrony which allows functioning systems of various kinds to be available at birth and at later stages of maturation. At the moment we have models at the simpler stage of studying morphogenesis of single organs - an apparently necessary way-station in the evolution of our models before we can tackle the synchrony problems of systemogenesis. At the moment, we look at one organ in a system and try to explain what sort of cellular interaction can give rise to its shaping. We may hope that, later on, when we understand this, we will have the intellectual apparatus in place to combine together our models of several systems to understand what sort of overall mechanisms allow coordination of their development.

Having established, in Section 1, the cellular level as an appropriate intermediate between the study of macromolecules and organelles by the biophysicist and biochemist, and the study of organismic control by the computer and information scientist and having now seen the interest of understanding embryological processes, let us briefly mention some of the mechanisms at the cellular level which shape the overall form of the organism, including that of its nervous system. (The reader will find an excellent overview of "the forces that shape the embryo" in Trinkaus [1969].)

One mechanism whereby a tissue may change its form is that of the autonomous change in cell shape. For example it is now well-known that various microstructures may be synthesised within cells during characteristic changes of shape, and that their destruction impairs such changes. Thus cells seem able to elongate themselves by producing microtubules aligned parallel to the axis of elongation. Again, cells seem able to constrict a portion of themselves by producing microfilaments which can then contract to provide the constriction by a sort of "purse-string effect" [Baker and Schroeder, 1967]. Schroeder [1970] has combined such mechanisms to provide an elegant model of neurulation - the process whereby a plate of cells on the back of the embryo is formed into a trough which then rolls up into a tube running the length of the embryo to then disappear beneath the surface of the back to form the rudiments of the spinal cord and brain.

Another mechanism whereby a tissue may change its form involves the combined effects of cellular adhesiveness and cellular motility. Such a mechanism helps us understand situations in which the attachments of cells change over time, but where there seem to be important specificities in the ensuing pattern of cellular attachments. Gustafson and Wolpert [1967 - for an exposition see also Wolpert and Gustafson, 1967] have given a masterly analysis of cellular movement and contact in sea urchin morphogenesis. Ede and Agerbak [1968] have been able to correlate changes in adhesiveness of cells (and the consequent change in their motility) in normal and talpid^3 mutant chick embryos with changes in the developing limb pattern in these embryos, while Ede and Law [1969] have expressed this correlation in the specific form of a computer simulation of limb development.

While elegantly showing how changes in cell shape, motility or adhesiveness can provide mechanisms for morphogenesis - both in the nervous system and elsewhere - the above schemes do not make explicit how a cell "knows" what contribution it is to make in the overall pattern. It is for this reason that other workers have developed the idea of "positional information". Here, the line of argument runs "If the cell is to change appropriately it must have information about its position within the embryo (and perhaps it will need to consult a clock, too)." An early approach to such positional information was in gradient theory (e.g. Child [1941]) - if a source of some metabolite were located at one end of the axis and a sink at the other, with a uniform gradient in between, then the concentration of metabolite in any cell would signal its position on the axis. Wolpert [1969] has suggested ways in which such a model needs refinement and elaboration, and Goodwin and Cohen [1969] have instantiated Wolpert's ideas in a model in which position is signalled by the phase differences between families of pulses propagating with different delays from cell to cell. By contrast, automata theorists have shown how cells may be formed into complex arrays without explicit "addressing". Rather, each cell is capable of a finite number of states, and at any time the cell changes state in a way dependent upon its previous state and that of its neighbors. For example, von Neumann [1966] exhibited a self-reproducing array with tens of thousands of components, but the cells were only capable of 29 states, and so could not "know where they were". Arbib [1967] has attempted to place this approach in a more biological context. The work of Apter [1966] should also be mentioned here. Other authors have compared the change of state rules used by von Neumann and others to the rewriting rules employed by linguists to "grow" a sentence from its grammatical description, and are now exploring the applicability of formal linguistics to theoretical embryology (Lindenmayer [1968], Laing [1969]).

In considering the specificity of cellular connections, we must not be misled by estimates that the amount of information in DNA is far less than that contained in the connections of the brain, which some have taken to imply that connections in the brain must be random. To see this, consider the following computer program which comprises four instructions:

1. Set n equal to zero.
2. Print out n.
3. Replace n by n + 1
4. Return to the second instruction.

If you observe a computer executing this program, it will emit a stream of numbers which is endless - at least till you have exhausted the capacity of the computer. Arguments that a comparison of the number of DNA bases with the number of connections in the brain shows that the brain must be a random network is as naive as comparing the four instructions of the above program with the number of positive integers and concluding that the sequence of positive integers, since it has more than 4 elements, must be a random sequence!

From programming computers we know the flexibility of programs having loops within them which are hierarchically structured to provide for a great deal of economy in the way we specify processes. As a biological example of a plausible "use" of such "nested subroutines", we may cite the retina of the frog. We have already cited the structure of the retinal output, but turning to the circuitry within the retina, we may note that Lettvin and Maturana have schematized the connections between the interneurons of the retina's second layer and the ganglion cells as falling into two or three segregated layers. A plausible wiring scheme would then prescribe that certain types of axons from the interneurons terminated in one layer and so are highly likely to connect one level of the dendrites of the ganglion cells while other types of axons bearing different transforms of the visual input would terminate in the other layer thus hitting other parts of the ganglion cell dendrites. By this means, one can very simply specify how to get a retina that would function perfectly for the frog trying to snap flies in its world, without having to specify point-by-point interconnections. Hence, a sort of "nested subroutine" approach could probably explain a great deal of the specificity of the nervous system without requiring an immense investment in genetic material. Such economy of genetic prescription augurs well for economy of functional description when we come to describe organizational principles for neurophysiological processes.

4. A DISTRIBUTED-COMPUTATION MODEL FOR THE CONTROL OF GAZE

Pitts and McCulloch [1947] 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 $\phi_o = T\phi$ [Arbib and Didday [1971] 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 $E(\phi) = 0$.

There are two main implementations of such a feedback scheme, only the second of which was considered by Pitts and McCulloch. (Henceforth, let us use W_ϕ to abbreviate $\mathcal{W}(E(\phi))$.):

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 <u>tracking scheme</u>, 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 <u>repeated</u> application of the error-correction the error eventually go to zero. A controller implementing tracking may proceed according to one of two schemes. The <u>implementation</u> mode corresponds to continually modifying the pattern until one is found which is in standard form; the <u>planning</u> mode corresponds to continually modifying the transform until one is found which will bring the given pattern to standard form:

I. Here ϕ will be the latest transformed version of the input pattern.
   1. Replace ϕ by the new input pattern.
   2. Use $E(\phi)$ to obtain W_ϕ

3. Form $W_\phi \cdot \phi$ to obtain the new pattern ϕ
4. Is the new $E(\phi)$ close enough to zero?

   YES: Exit, ϕ may be treated as in standard form

   NO: Go to 2

II. Here ϕ will be the fixed input pattern, and T will be the updated transform to be applied to ϕ

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. [Arbib, 1972].

The hard work in such a scheme is in 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 [1945, 1946] 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 (Pitts & McCulloch, [1947], 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. 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.

5. DISTRIBUTED MOTOR CONTROL

It should be noted that even if the mathematical equations formalizing the Pitts-McCulloch scheme 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 [1960] 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 [1972], 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 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.

The controller then elicits a ballistic movement by firing three 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 a distributed motor controller of a type 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.

6. REDUNDANCY OF POTENTIAL COMMAND

To exemplify this discussion, let us present a model of the reticular formation, and then relate it to a model of frog visumotor behaviour which involves layered distributed computation. 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 the gaining of access to a program rather than the execution of the 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 [1958] 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 eachwith 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 [1968] designed and simulated the compartment model, called S-RETIC, 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 4) 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 <u>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</u>.

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 neurons 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 [1971 a,b] 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 has an array of modules which must interact to get a majority favoring the same mode, the task of the second layer of Didday's hypothetical tectum 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. [1959]. 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.

7. ORGANIZATIONAL PRINCIPLES

To place our discussion in perspective, let us make explicit the effect our views of brain structure will have upon our approach to modelling brain function, contrasting three possible strategies for theoretical neurophysiology with the principles that guide our own approach.

In certain invertebrates, we may find that the function of the system we wish to explain is mediated by a rather small neural network and so we might actually hope to track down, by explicitly simulating the behaviour of say 100 or so neurons, all the details of their interaction, and so obtain a plausibly complete explanation of how a locust, say, walks or flies. (See, for example, the beautiful review of 'Insect Walking' by Donald Wilson [1966]).

When we turn to vertebrates, this strategy does not work, save in studies of peripheral circuits for muscle control, for there are just too many neurons. There are various strategies to take, depending on our ideas about structure, as to how one might make a model. The physicist has one ready answer for how we might model a system with millions or even billions of neurons. From his study of gases he would suggest statistical mechanics [or - in technical terms - "average the hell out of it".] Unfortunately, such averaging may destroy the very parameters of interest to us if we want to explain linguistic behavior or coordinated motor behavior as in a frog snapping at a fly. On the other hand, if we want to understand how the cooperative behavior of many billions of neurons in the cortex gives rise to evoked potentials or electroencephalograms, then some sort of statistical mechanical approach may well be worth while. However, a straightforward statistical approach to the very large system will not do for more detailed structural questions about complex information processing in brains.

In this context, it may be worthwhile to contrast two types of statistical model. Winograd and Cowan [1963] were concerned with the fact that in as large a system as the brain one has both the likelihood of not specifying completely all neural interconnections accurately by genetic parameters and also the likelihood of many "malfunctions" of components during actual information processing by the system. They wished to design networks with enough redundancy to insure that the organism would not be too unreliable. Their strategy was to start from a very specific function they wanted a hypothetical "nervous system" to undertake, and then provide ways in which they could transform this "nervous system" into a new form which was sufficiently redundant that quite a lot of sloppiness in the "wiring" and in the behaviour of the "neurons" would still give correct overall function. This strategy of starting from a specific structure for computing some function and finding ways of introducing redundancy to make it resistant to certain types of damage both in growth and function is radically different from the strategy Cowan [1969] took in his later work, in which he looked at interactions between thalamus and cortex only in terms of gross statistical parameters of their interconnection, and then asked if certain aspects such as cortical rhythms could be explained on this basis. In this case one only wants some crude parameters of overall system behaviour such as the period of rhythm recorded in gross potentials, and so one can "average out" a lot of detail by statistical mechanical techniques. But if one wants to look at the detailed state-dependent processing of inputs to get outputs then one has to impose far more structure, and study deterministic operation at a certain level.

Another approach to modelling a large system is that of compartment models. A brain modeller taking such an approach will not try to average over the complete system, but will look at the gross anatomy of the brain to subdivide the brain into various regions. He will thus try to simplify the problems of explaining one large region of the brain by breaking it down into a collection of interconnected "black boxes" and see if by making multiple plausible guesses about those boxes and their interconnections he can put together a reasonably functional model of the overall system. It may then be easier to take those individual boxes with their plausible functions and try to model them back down to the cellular level than trying to do the whole thing directly. Perhaps one of the most interesting brain models of this kind is that of Kilmer, McCulloch and Blum [1968] on the reticular formation. As we saw in the previous section, they used neuroanatomy to legitimise the compartmentation of the RF into a series of "modules" ascending the longitudinal axis. Each module could then be modeled as a whole, and then the simulacra could be interconnected to get the overall change-of-mode behavior which they posit to be exhibited by the reticular formation.

With this as background, I want to suggest eight principles which may help us understand how the human brain can control the complexities of a human's behaviour. [Their elaboration appears in my book "The Metaphorical Brain"]:

1. <u>Theory Must be Action-Oriented</u>:

One often talks as if human perception merely involved being able, when shown an object, to respond by naming it correctly. However, it is often more appropriate to say of an animal that it perceives its environment to the extent that it can interact appropriately with that environment. We can perceive a cat by naming it, true, but our perception may involve no conscious awareness of its being a cat per se, as when it jumps on our lap while we are reading and we simply classify it by the action we take as "something-to-be-stroked" or "something-to-be-pushed-off". In computer jargon, then, we may say 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.

2. <u>Redundancy of Potential Command</u>:

To repeat the argument of Section 6 - you perceive something and yet may 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?" In other words a brain must be able to "Resolve Redundancy of Potential Command."

We next state our third principle which, as we commented in Section 2, is known to all neuroanatomists but has been strangely neglected in brain theory:

3. The Brain "is" a Layered Somatotopic Computer:

In relating such a principle to statistical brain theories, I would suggest that we look at the brain in terms of specific structures, and only use randomization when we really feel we are ignorant. In other words, it is a justified strategy to make a model in which only certain parameters can be confidently specified - either by experiment or previous theory - and thus to set up random values for all the other parameters in the model. One may hopefully explain some of the functions for the system being modeled despite our gap of knowledge about structure, and then gain more detailed understanding of other functions as we find out more about further details of structure. This is something of an evolutionary strategy of biological modelling. What our third principle says 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. We must add that muscles are not unitary devices with a single controller, but instead each muscle comprises a multitude of fibres controlled by a whole population of neurons. Since an _array_ of input cells must thus control, with the aid of internal variables, an _array_ of output cells, it seems at least plausible that it is an _array_, rather than a central executive, that intervenes. We thus have made the hypothesis, encouraged by Sections 4 and 5, that, unlike most digital computers,

4. The Brain is a Distributed, Highly Parallel Computer

Our next point has been well introduced by our discussion of Section 3, which we may augment by suggesting that certain vagaries of brain structure only make sense when we see that the human brain is a variation on an evolutionary theme - and herein lies much of the power of _comparative studies_ which use experiments on animal brains to help unravel the knot of human mentation. Similarly, we must note that the brain had to grow, rather than being wired-up by a technician who could refer item by item to a blueprint. Thus in explaining organismic behaviour we must stress that:

5. Brains have Evolved; Further Each Brain has Grown:

Evolution has given each animal a basic repertoire of skills for survival. The frog brain enables the frog to snap at flies; the human brain enables the human newborn to suck and gaze and breathe and excrete. But in all animals to some extent, more so in mammals and perhaps exceptionally so in man, these evolutionary skills are augmented by individually acquired skills and memories. The social basis for much of human skills is, in fact, so great that nowadays man's adaptation of and to his environment is more dramatically a process of cultural than of biological evolution.

These experiences, skills, memories cohere into what has been called an internal model of the world. This model is not a cardboard replica, but rather the memory structure that, for example, lets us walk into a strange room and, on the basis of visual stimuli from a brownish rhombus, know that a table is present and that we may put the papers we are carrying on that table without risk that they will fall to the floor. Thus:

6. The Brain must be able to Correlate Sensory Data and Actions in such a Way as to Build up an Internal Model of the World.

Of course it is not enough to perceive the presence of a table on which we place our papers if we then release our grasp of them some three feet from the table. Too many discussions of perception overlook our first point (action orientation) and so talk as if it were enough to classify an object. But we must know where it is if we are to interact with it:

7. Perception is of Where as well as of What

Finally, we note that when walking around obstacles, we decide where to walk but, unless the ground is very uneven, need not concentrate on placing our footsteps; further, it appears that the midbrain control of stepping can relegate to spinal mechanisms the maintenance of an upright posture as we step. These, and other, considerations (see also Greene [1964, 1968, 1970]) suggest that it is useful to theorize that

8. Both Structurally and Functionally, the Brain is Hierarchically Organized.

As we saw in Section 6, the study of frog behavior - which is very much in the spirit of these principles (though simplified by the non-adaptive aspect of the behaviour modelled) - sheds new insight on RF modelling, and suggests alternate hypotheses. Though the model is still but a crude oversimplification of the complexities of a real frog brain, we believe that our partial successes show that these organizational principles (or their evolved descendants!) will play a crucial role in future theoretical neurophysiology.

REFERENCES

P. K. Anokhin (1964):Systemogenesis as a General Regulator of Brain Development, in 'The Developing Brain' (W. A. and H. E. Himwich, Eds.) Progress in Brain Research 9, 54-86.

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

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

M. J. Apter (1966): Cybernetics and Development, Oxford: Pergamon Press.

M. A. Arbib (1967): Automata Theory and Development, I, J. Theoret. Biol., 14, 131-156.

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, 3-18. (Part II, "The Hologram Metaphor," by M. A. Arbib and P. Dev, will appear in a later issue.)

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

P. C. Baker and T. E. Schroeder (1967): Cytoplasmic Filaments and Morphogenetic Movement in the Amphibian Neural Tube, Developmental Biology, 15, 432-450.

C. C. Boylls and M. A. Arbib (1972): Organizational Principles in Cerebellar Modelling, Progress in Biophysics, (to appear).

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

C. M. Child (1941): Patterns and Problems of Development, Chicago: Univ. of Chicago Press.

J. D. Cowan (1969): Statistical Mechanics of Nervous Activity, Technical Report JDC 69-1, Committee on Mathematical Biology, University of Chicago.

R. L. Didday (1971 a): Simulating Distributed Computation in Nervous Systems, Int'l. J. of Man-Machine Studies, 3, 99-126.

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.

D. A. Ede and G. S. Agerbak (1968): Cell Adhesion and Movement in Relation to the Developing Limb Pattern in Normal and Talpid^3 Mutant Chick Embryos, J. Embryol. Exp. Morph. 20, 81-100.

D. A. Ede and J. T. Law (1969): Computer Simulation of Vertebrate Limb Morphogenesis, Nature, 221, 244-248.

B. C. Goodwin and M.H. Cohen (1969): A Phase-Shift Model for the Spatial and Temporal Organization of Developing Systems, *J. Theoret. Biol.*, *25*, 49.

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

P. H. Greene (1967): Models for Perception and Action, *Proceedings of the First Annual Princeton Conference on Information Sciences and 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."

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.

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]).

T. Gustafson and L. Wolpert (1967): Cellular Movement and Contact in Sea Urchin Morphogenesis, *Biological Reviews*, *42*, 442-498.

M. Jacobson (1970): *Developmental Neurobiology*, New York: Holt, Rinehart and Winston.

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. Mesarović, Ed.), New York: Springer-Verlag, 286-375.

R. Laing (1969): *Formalisms for Living Systems*, University of Michigan Report 08226-8-T.

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.

A. Lindenmayer (1968): Mathematical Models for Cellular Interactions in Development, *J. Theoret. Biol.* *18*, 280-299, 300-315.

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.

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.

T. E. Schroeder (1970): Neurulation in Xenopus Laevis. An Analysis and Model Based Upon Light and Electron Microscopy, *J. Embryol. Exp. Morphol.*, *23*, 427-462.

J. P. Trinkaus (1969): Cells Into Organs: The Forces that Shape the Embryo, Englewood-Cliffs, N.J.: Prentice-Hall.

J. von Neumann (1966): Theory of Self-Reproducing Automata, U. of Illinois Press (edited and completed by A. W. Burks).

P. A. Weiss (1941): Self-Differentiation of the Basic Patterns of Coordination, Comp. Psychol. Monogr. 17, 1-96.

D. M. Wilson (1966): Insect Walking, Ann. Rev. Entom. 11, 103-122.

S. Winograd and J. D. Cowan (1963): Reliable Computation in the Presence of Noise, Cambridge, Mass. The MIT Press.

L. Wolpert (1969): Positional Information and the Spatial Pattern of Cellular Differentiation, J. Theoret. Biol. 25, 1-47.

L. Wolpert and T. Gustafson (1967): Cell Movement and Cell Contact in Sea Urchin Morphogenesis, Endeavour, 26, No. 98 (May) 85-90.