402

AF61(052)402

AD633002

SYSTEM RESEARCH LTD

1

THE CYBERNETICS OF EVOLUTIONARY PROCESSES AND

OF SELF ORGANISING SYSTEMS

BY

GORDON PASK

20 HILL RISE
RICHMOND
SURREY
ENGLAND
RIC 0801

THE CYBERNETICS OF EVOLUTIONARY PROCESSES AND OF SELF ORGANISING SYSTEMS

By

GORDON PASK.
(System Research Ltd)

The text of an invited Paper delivered at the 3rd Congress of the International Association of Cybernetics, Namur, (Belgium), in September 1961.

I would like to thank the research sponsors who lend substance to these ideas by giving us the chance to apply them in our work, in particular, to Burroughs Research Corporation (Cognition Project), Solartron and Wright Air Development Division (through Contract AF 61, 052,402, on Group Teaching Systems).

Gordon Pask.

The Cybernetics of Evolutionary Processes and of Self Organising Systems

1.1 Introductory Comment

Cybernetics is an art, a technology, a science and a philosopy. Neither the art of effecting control nor the technique of applying this art, is my speciality. But I intend to speak as a Scientist and later as a Philosopher about these other aspects of our discipline.

Cybernetics is the science of systems. To-day we shall consider a couple of special but important categories of systems, namely, those that are self organising and those that evolve.

First let us be certain what a system is. Perhaps you will excuse a lengthy and rather elementary discussion on this point because my contentions can readily be mis-construed, and to avoid this we must be quite certain of the meaning given to a word that is so often and so carelessly bandied around, as the word 'system'.

1.2 Definition of a System.

The paradigm of a system is Ashby's(1) concept of a black box. The idea is closely related to Shannon's statistical information theory, with which, indeed, it grew up. Within the compass of Shannon's theory, a black box is an information source. Within a more general context, the contents of a black box may be anything you please. The black box may, figuratively, be a radio set with components inside, or it may be this room with you and I inside it. Of course, we do not mean that the walls of the radio set or this room literally bound the box, for the intended closure is informational. The whole construction is meant to emphasise that certain relevant attributes have been decided, and that the set of these determines a state description, for the values of these attributes, at a given instant, are an assertion of the state of the black box at that instant.

In one plausible state description we might consider all the people in this room and where they are sitting, who is talking, and the temperature of the room. We neglect, as irrelevant in this state description, the concentration of tobacco smoke, and the intensity of the illumination. So a state is specified by your positions, the temperature, and the fact that I talk continually. But it is essential to remember that for every attribute of the physical assembly that is deemed relevant, there are an indefinite number of equally acceptable attributes that are deemed irrelevant.

Ashby comes to the crux of the matter by insisting that we talk of systems, not of things.

The attributes may be independent binary properties that are or are not present, such as I am, or am not, talking. If they were all like this the state could be expressed as a binary number, and for n attributes there would be 2_n states. But there may be constraints, for example, the fact that not more than one of you can sit on a chair at once, is a constraint and the fact that the temperature cannot rise from 22°C to 24°C without going through 23°C, is also a constraint. Consequently only certain arrangements of you need be considered, and temperature ceases to be a binary attribute of the form 'temperature is or is not 22°C', 'temperature is or is not 23°C', and becomes simply 'temperature', which assumes a lot of different values.

We suppose there are measuring instruments that convert attributes of the real world into abstract symbols, often selected from the set of numbers. Typical measuring instruments are voice keys, and thermometers. An attribute is defined by the construction of the instrument that measures it. Equally, however, these instruments may be our own senses and in this case the specification of an instrument is called a 'percept'. Now, whether we do the job with our own senses, or whether we allow an instrument manufacturer to do it for us, we end up with an 'abstract representation' or merely 'an abstraction' of the real world, which consists of a set of possible numbers, defining states in a state description determined by the choice of attributes, and conveniently envisaged as the readings obtained from a set of meters or dials on the outside of the black box. The meters or dials are labelled as 'variables' with which the attributes have been identified by the process of abstraction. Thus, as before, the values of the variables define a state of the black box.

The totality of possible assertions about the variables and their relations to one another is called 'a universe of discourse'. Statements that are understandable within the universe of discourse are made in terms of an observational language. Some conceivable statements will be excluded by any constraints that are agreed to exist by a group of observers who communicate with each other in this observational language. That temperature is a continuous variable, is taken as logically true in the universe of discourse we have thought about so far, and consequently an assertion like, 'the temperature went from 22°C instantaneously to 24°C', is logically false. Similarly, 'more than one of you is sitting on the same chair' is logically false. There are only some logically possible assertions.

A system is a universe of discourse together with its identification to the attributes in a state description of the real world. States of the system are sets of values of these identified variables.

Notice that any assertion about the fact that the variables are identified with the real world, of the room, or you and I, is prohibited in the systems observational language. Such assertions could only be expressed in a metalanguage used for discussion about the system as such and its relation to other systems.

1.3. Dynamic Systems and Behaviour

We shall be wholly concerned with dynamic systems. In the present case this means that the room containing you and I is freely supplied with energy, mostly derived from metabolising our breakfast, so that we produce a lot of autonomous activity as a result of which the system describing us changes state and provides whoever constructed it (call him the 'Observer' or the 'Experimenter') with a stream of evidence about the attributes with which his system is identified. This stream of evidence is a behaviour of the system. A behaviour delineates those events that actually do occur (as distinct from those that are logically possible). At this point, we come up against a basic uncertainty in measurement, a so-called metrical uncertainty which gives cogency to the idea of the black box. There is a definite limit as to how often or how accurately the values of the variables can be specified. In terms of a black box, the meter readings or dials indicate uncertain evidence regarding the values of the attributes, which they are supposed to measure. The system defines the limits and the location of an observer's certainty.

Recall that in informational thinking, the black box is an information source. Similarly an Observer becomes a receiver aware of the possible states of a transmitter connected to the source. The universe of discourse is a specification of just these states and the constraints upon their transition. The abstraction specifies the coding that determines the form of channel, namely the connection of the transmitter and its relation to the receiver. A transmitted message constitutes evidence. A sequence of received messages constitute a behaviour. If we pose a source of 'noise' acting into the channel, this 'noise' represents the Observer's metrical uncertainty. Like any other model of this kind it is not a picture of things as seen by the Observer himself, but a picture as seen by someone looking on from outside at the process of observation. In other words, this model of observation exists in a metalanguage. With this in mind, the image in DIAGRAM 1 is concise and convenient.

Now, an Observer would like to describe a coherent behaviour in terms of his observational language, so that he can make predictions about the state of the system. The possible predictive hypotheses depend upon the universe of discourse chosen by this observer, or, equivalently, upon the constraints implicit in DIAGRAM 1. Predictions are hypotheses of which the observer has become more or less certain because of the evidence he has received, but because of his necessary uncertainty there is no more than a definite probability that a predicted occurrence will be realised.

The simplest procedure an Observer can adopt for collecting his evidence is cumulative induction. He watches what occurs, records how often it occurs, and represents this counting of events in terms of a probability, of occurrence, or joint, or successive, or conditional occurrence. Now he assumes that events which have occurred often will be likely to occur again. But this never yields certainty. For something can occur that has not previously occurred. Indeed, cumulative induction is rather inefficient and its great advantage is that it can always be used, even if nothing is known about the constraints that determine the behaviour of a Black Box. Whenever something is known, so that there are a limited number of logically possible hypotheses, the experimental procedure of science is much more informative. Using it, the experimenter poses an hypothesis and uses evidence in an attempt to disprove it for whereas indefinite repetition cannot entirely confirm any hypotheses, one negative observation can disprove a sufficiently strong hypothesis.

The system is greatly enhanced by adding a further channel as in DIAGRAM 2, using which an observer is able to change the conditions in the 'Black Box' or the corresponding 'parameters' of his system. In this case observations are made conditional upon different parameter values. The experimenter, of DIAGRAM 2, being in a position to control the parameters, or at least to control their probable values, is able to perform experiments upon the system, rather than watching and waiting for events to occur and the falsifying procedure becomes realistic. A parameter in the system corresponds with some action upon the real world that is correlated with a coherent change in attribute values. Thus the action of opening the window corresponds to a parameter of the system we have been examining, so far.

A system is a structural framework largely determined by an observer and the observer cannot guarantee before he begins his experiments, that it will exhibit coherent behaviour. If it does not, it is uninformative and the observer must try another. But there is an indefinitely large set of possible systems and no real observer can search through the whole set unaided. In practice, his search, if it is realistic at all, must be guided by previous experience - (which is only useful in an ordered world) - or by convention - (which is only tenable - in an ordered world).

But suppose the system is coherent, it can be variously characterised. In particular, we are interested in statistical measures like the variety of its behaviour and in measures of the organisation it exhibits. If the system is coherent, it will be possible to compute a measure of organisation as a redundancy, which is:

$$\text{Redundancy} = \eta = 1 - \frac{\text{Variety}}{\text{Maximum Variety}} = 1 - \frac{\xi}{\mu}$$

and having computed a redundancy to ask what one would mean by a self organising system.

1.4. A Self Organising System

The most rigorous discussion of a self organising system is due to Beer(3). But for this purpose, we shall adopt a very elegant definition advanced by von Foerster(4) which is consonant with Beer's argument. Von Foerster points out that a system is a self organising system if and only if the rate of change of redundancy is positive valued:

$$\frac{d\eta}{dt} > 0$$

Further, von Foerster argues that this system can be self organising either because the variety of the behaviour is reduced or the maximum possible variety is increased.

If

$$\mu = \text{Constant}, \quad 0 > \frac{d\xi}{dt}$$

If

$$\eta = \text{Constant}, \quad \frac{d\mu}{dt} > 0$$

Suppose that μ the maximum variety is constant, as it would be in a conditional probability machine, like one of Dr. Uttley's devices or a finite adaptive network. Statistical constraints are built up in the network when its input is manipulated and the redundancy of its behaviour increases. This is a self organising system which exists over a definite interval, but ultimately when all the constraints are built up, $0 = \frac{d\xi}{dt}$.

Suppose, on the other hand, that the variety ξ, is made constant, and the maximum variety is increased by adding elements, for example, by uniform growth, as in a crystal, or in the development of a mushroom, where after a primordium has appeared, the mushroom grows uniformly. This is also a self organising system, but unless we are prepared to countenance an indefinite extention, it is liable to a different kind of limitation.

In the commonest and most interesting cases there is considerable interaction between these modes of organisation and development and the process as a whole is called 'differentiated growth'. A structure develops, in the simplest case by addictive growth alone. Then, within the limits of this structure the degree of organisation will increase, after the manner of the constraints in an adaptive network. The point is that we cannot, in these conditions, say what is becoming organised for whatever it is (the structure) is, itself, evolving. Looking a little further back (at the system we considered a moment ago), the primordium is a structure which evolves from an increased organisation amongst the elements of a completely different structure, namely, a mesh of hyphae. The apparent discontinuity between hyphae and the primordium is, of course, typical of biological development. Comparable differences appear between cellular level and tissue level organisation and between individual and group organisation.

1.5. The Strict Impossibility of a Self Organising System

We defined a system as an identified universe of discourse. We defined a self organising system as a system such that

$$\frac{d\eta}{dt} > 0$$

Remarking that its continued existence depended upon a change in the value of μ. But, whereas ξ, the variety, is a function of the behaviour of a system μ depends upon the state description, and the specification of the system itself. Further, μ cannot change without change in the system. Thus, a strictly interpreted self organising system is either trivially restricted (μ remains invariant, ξ changes) or it is a logical impossibility (for change in μ changes the system we are talking about). All the same the idea of a self organising system is intuitively reasonable and seems to ideally describe the observational framework in which we are accustomed to image events like the development of organisms. So, rather than discarding the notion, it is worth elaborating the underlying informational construct to permit its rigorous definition. Minimally, we must replace the idea of selective information by the more comprehensive theory of 'Scientific information' advanced by Gabor and MacKay.(5) Whereas in Shannon's concept, a source, transmitter and channel, must be decided before the measure 'Selective Information' is introduced and the whole notion depends upon a satisfactory specification of these structural entities, the Gabor and MacKay theory is also concerned with the act of choosing a structure which is tantamount to choosing what can be asserted in addition to what does occur. In our nomenclature the structural aspect of the theory is concerned with the universe of discourse and its identification, or, in other words, with the metalinguistic specification of a system, whereas its metrical aspect is concerned with the acquisition of evidence given these structural constraints (if a system is defined the theories are isomorphic). An observer is liable to a structural uncertainty (about the form of hypothesis to test or enquiry to make) in addition to the metrical uncertainty which we considered previously.

1.6. Structural Uncertainty.

Thus it is useful to distinguish between (i) the case in which an observer has no significant structural uncertainty and the Black Box analogy is completely applicable. (ii) The case in which structural uncertainty is important.(6)(7)(8) In either case it will be possible to construct a system. But whereas in case (i), due to the assured relevance of the identified attributes, the system will be informative, the system of case (ii) is only informative by dint of good fortune, experience, or perhaps intuition.

Typically, (i) is characteristic of the classical sciences where there are well established methods of measurement (the construction of a thermometer or a voltmeter is public knowledge) and where measurements are comparable (a thermometer and a voltmeter although different instruments, can be compared within the boundaries of classical science). Further, all tenable hypotheses appear in a common frame of reference and data gleaned with respect to one hypothesis tends to validate or deny all others. In short, any system is a subsystem of a gigantic system that defines the entire discipline. The limit of case (i) is reached when a designer observes his own automaton. There is no structural uncertainty for the function of each part is known.

In contrast (ii) occurs in the Behavioural Sciences where there is no common frame of reference and knowledge is only locally consistent. Often the experimenter does not know, before he performs his experiment, what kind of enquiry will be relevant and will yield a coherent behaviour. Thus, a child psychologist is bound to give his subject a great deal of liberty. He can 'motivate' the subject and he can 'direct' his attention, but, in the last resort these are chancy procedures and the experimenter's system (if it is to be informative) must be identified with whatever does occupy the subject's attention. Building blocks at one moment, toy motor cars, at the next. Sometimes, of course, the experimenter does know what kinds of enquiry will be needed, but, in this case, the different enquiries tend to be incomparable (and thus they cannot necessarily be identified with variables in the same system). Investigating the development of embryos it is necessary to perform experiments which refer to incomparable entities such as organisation at a cellular level and the organisation of a tissue and the embryologist is likely to obtain incoherent results if he adheres to only one approach.

1.7. Application of the Phrase 'A Self Organising System'

Now for Case (i) the phrase 'A self organising system' is unnecessary and (as we have seen) the idea it conveys is inconsistent and apt to be confusing. For, if there is no structural uncertainty, an Observer who obtains incoherent results must either conclude that the world is incoherent (and thus behaves like a chance machine) or that his method is inadequate. If the latter is true, there exist procedures for improving the method which entail making more accurate or more prolonged observation, and which amount to changes of coding of the information from a given Black Box.

Consider a world in which behaviour is manifest by the appearance of a pair of differently coloured (Red and Orange) signal lamps which are placed close to each other on a panel and labelled as X and Y. Suppose that the most accurate state description we need to consider is identified with the condition of these signal lamps according to the mapping rule State a ≡ X(R), State b ≡ X(C), State c ≡ Y(R) and State d Y(O). When considered in this detail an observer is presented with a system (I) having a determinate behaviour such as:

(I) 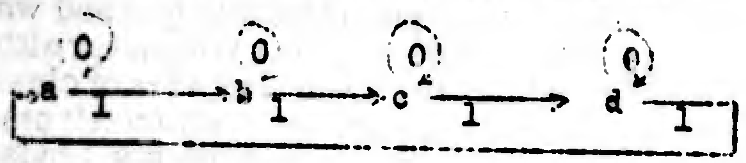

In the state description, let A = [a or b], B = [d or c], and let D = [a or c] and C = [b or d]. An observer who views the panel through a red glass filter (thus changing his identification for the colours become indistinct) considers a system (II) with a behaviour:

(II) 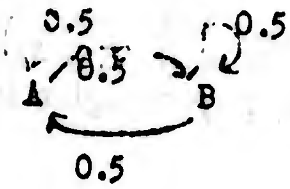

whilst an Observer who views it through blurring spectacles (so that the position is indistinct) will obtain a system (III) with a behaviour that preserves the essential feature of a cyclic transformation in (I), but which requires no greater observational channel capacity than system (II) which obscured this characteristic.

(III)

Finally, let A(A) mean 'state A after state A', and let A(B) mean 'State A after state B', so that an Observer, viewing the world through a red glass filter (but able to recall previous states) discerns a determinate behaviour (IV) which is isomorphic to the determinate behaviour of (I) given the correspondence: A(A)→b, A(B)→a, B(A)→c, B(B)→d, namely:

(IV)

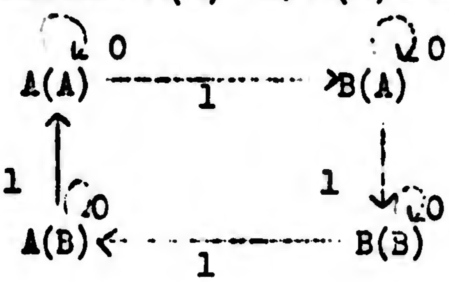

Now, an Observer viewing the world through a filter discovers a behaviour as in (II), apparently determined by a chance machine. He can make no predictions about the state of his system. The procedures he can use to improve his ability in this respect are:

(1) Increase the accuracy of observation by removing the filter to obtain (I). (2). Prolong the observation and recall previous states to obtain (IV), or finally, change the state description without altering the channel capacity by exchanging the filter for the blurring spectacles to derive (III). If this is a Case (1) system, the coloured filter and the blurring spectacles are assumed to be comparable impediments defining he deficiencies of comparable procedures of measurement. In particular, the experimenter is able to explain what he is doing when he changes the coloured filter or the blurring spectacles in much the same way that he can explain the mechanism of sensory instruments and it is because of this that the procedure is admissible.

For case (ii) there are several possibilities. Notice, in the first place, that an apparently haphazard world cannot be brought to heel by the neat 'coding' procedures we have just described and the observer is bound to search for a form which on the one hand seems regular (which is one of his percepts) and in terms of which the behaviour of the world is coherently representable. He has to construct a completely different Black Box. We remarked in 1.3. that this search (for a system) could not be unlimited. There are indefinitely many attributes of the world, but either the observer or the observed entity will have decayed before the list is exhausted.

Failure to discover some kind of apparent regularity means that the observer is confronted with events of which he is unaware either because he cannot conceive them, or because they cannot be expressed, for communication, in the language of his group.

We might, with caution, call the abortive system that our observer is trying to construct, a 'haphazard system' (which is quite distinct from 'chance machine' behaviour within a well defined system of possible outcomes), since experience of it could vary from 'nothing' to fleeting, elusive, glimpses of regularity. (This sort of world cannot, so far as I know, be 'found' but it is non trivial since its conditions can be approximated, over a limited interval, by a special arrangement. Take a display presenting varied and possibly regular occurrences. Consider an indication, that an observer, viewing the display, detects any kind of regularity. Let X be the occurrence on a given occasion. If the observer detects regularity, decrease the probability of X upon succeeding occasions, either manually or by the action of an automatic device).

At the other extreme, the Observer finds a system in which an interesting body of behaviour is always coherent. Thus Lettvin and Matturana(9) discovered a rather elaborate experimental system in which the neural activity as well as the overt response of a frog to visual stimulation became coherent (the crux of their discovery was the set of logically elaborate retinal configurations that act as stimuli which are, 'simple' to the frog). Once the system is found, Case (ii) becomes Case (i).

Finally, the world may be such that an observer who is anxious to find coherence must construct a self organising system which now appears as a sequence of systems related by agreement on the part of a group of observers but (because they are commonly incomparable), unrelated in any other way.

Consider the topic of intellect and its developments in the child. Each of Piaget's(10) 'stages' in the development presents a coherent behaviour (the descriptive invariants of this behaviour characterise the 'stage'). But each is observable in quite a different experimental system, and, in order to give a coherent account of maturation (either of brain processes or intellectual behaviour) an observer is bound to construct a sequence of systems. Strictly these are incomparable and cannot be regarded as subsystems of a common whole. But these systems are related and called a self organising system because the child is something which the observer regards as a consistent entity for reasons that are nothing to do with his experimental data.

1.8. The Salient Characteristics of 'A Self Organising System'.

(1) The phrase is used relative to the existing state of knowledge. Previously incomparable attributes and systems become comparable as knowledge increases and mere correlation between events gives place to a substantial predictive rule. Given any particular case, we do not deny that it will become representable strictly in terms of a 'Black Box', nor that all observables will ultimately have this form (though if they did, 'scientific enquiry' would amount to statistical prediction). The idea of a self organising system belongs to the present and imperfect search for coherence.

(2) A self organising system is specified relative to a given observer or a group of observers with a common logic. We stressed in 1.3. that any system is defined relative to some observer, but this fact can be overlooked when, as in Case (1) the relation is determined by scientific convention.

(3) It is a sequence of otherwise incomparable systems, exhibiting coherent behaviour, rendered comparable by agreement and resort to such additional data as 'the behaviour is engendered by that man', or 'by the development of that species'.

(4) It is specified relative to a topic that interests the observer or relative to an objective to be achieved. The occurrence of self organising systems in the behavioural sciences is partly due to our ignorance in this field, but partly due to the kind of enquiry (regarding issues like 'development' or 'maturation') that is significant in this field.

(5) We have already remarked that an unguided search for a system would be impossible and, in practice, the scientist is directed by metainformation. Often this assumes the form suggested in (3). Other forms entail inferences of 'similarity', for example the observer infers that another man will 'learn' (and his behaviour may be represented as a self organising system because of this) on the grounds that this other man is made as he is made and of the same material (which has absolutely no bearing upon the observational language statements about information processing). Again the observer may infer that 'it' where 'It' perhaps is an automaton, is similar to himself in the sense that it will construct the same kind of concepts. Finally, we may cite as a commonly used bit of metainformation, our belief in the regularity and continuity of the environment.

2.1. Evolutionary Systems

When we are concerned with the evolution of a species or the evolutionary process which is the mechanical correlate of the behaviour called 'learning', our attention is directed towards the necessarily time dependent character of a self organising system. By definition we are talking about a time sequence of systems that becomes a single self organising system by virtue of the topic we have in mind and it is convenient to call such sequences, Evolutionary Systems.

Once there existed a rough-and-ready distinction between the physical and the biological science - the physical scientist sat in his laboratory and constructed artifacts or experimental situations which he devised at will and studied - the biological scientist went out and brought into the laboratory animals and vegetables which had been generated by evolutionary processes beyond his jurisdiction. Now when, as Cyberneticians, we are chiefly concerned with organisation this distinction disappears. We can perfectly well construct evolutionary processes within the laboratory that will generate any number of organisations from which we select a few that happen to be interesting. It is true that in certain essential particulars, the evolutionary process remains beyond our jurisdiction so that, even if it takes place in a computor, we are bound to learn about it by studying its products. The process itself is incompletely determined and our abstraction of the process is a typical Evolutionary System.

One objective is to discover correspondences between this evolutionary system and similar abstractions from the evolutionary processes of nature in an attempt to set out principles common to any kind of evolution. Consequently the Cybernetician may identify a collection of elements in his evolutionary system with the object called a population in the natural system insofar as the 'collection' and the 'population' exhibit the same behaviour. Also he will alter the evolutionary rules in his process (which in turn effects the development in his system), and try to increase the number of possible identifications. Finally, we are anxious to investigate modes of evolution that do not appear (or are not recognisable) in the real world because of its topology or because of energetic constraints. Rashevsky's(11) work in this field is well known and chiefly related to the evolution of species. Richard Feldman(12) has recently developed a process for the evolution of structures in networks that has an obvious correlate in the maturation and learning activity of a brain. I could do justice to neither in a brief summary, and refer you to the original papers. For the moment, I shall present an outline of a model of my own which, though less elegant than either Rashevskys or Feldmans model, is especially tailored to demonstrate the points I wish to make.

2.2. Background for Abstract Models.

Any model for evolution must involve the environment in which evolution occurs as well as that which evolves. If we are thinking of organisms that evolve, or species of organism, the environment will be a physical structure in which there is food to consume and form to perceive and the companionship of other organisms. If we are thinking of brain activities that evolve the environment will be sequences of messages in various languages and various modalities. These environments, or, for that matter, many others, can be represented as abstract constraints. One feature of the present model is that its abstract constraints can be identified with a variety of different environments or the attributes of different evolving entities. But these interpretations will appear more or less plausible (to the extent that some may become impractical and others become imperative) at different stages in the evolutionary process. (We might predict as much, since the evolutionary system is a self organising system). As a consequence we have a structural uncertainty regarding what it is that does evolve - the organism, an aggregate of organisms - or the process of development of each individual. Let us specify a few prerequisites for evolution of the kind we are trying to model:

(1) In the real world, evolution occurs when there are a number of distinct elements each of which can survive in certain conditions of the environment and cannot survive in others. The issue of whether or not a particular element does survive in conditions that permit its survival depends upon its behaviour, and in any interesting process the behaviour of the elements must be such that they tend to remain in existence. Survival of the physical material that constitutes the element is a prerequisite of the stability of the organisation that maintains the element, commonly by resynthesis of structural components from raw material (or 'food') in the environment. Thus, conversely, survival depends upon stability. Cells are typical elements in the biological environment (though we could cite the reaction centres of some autocatalytic reactions or regions of activity in a network of artificial neurones as perfectly legitimate evolving elements in different environments). Thinking of cells, the survival of the energy transforming mechanism is a prerequisite for maintaining the nucleic acids, that chiefly determine the cellular organisation.

(2) We remarked that survival is conditional. The simplest conditionality occurs if 'food' is available in short supply. 'Food' of course, can be read as 'money' or 'electric current' or any other conservable commodity without altering the essential condition that if it is restricted there will be competition between the elements for whatever is available.

(3) Either the elements must be capable of reproduction on their own account, given success in the 'Food Competition' or there must be a locally specified state of the environment such that one element is

created when the local 'Food Concentration' is high. Since it is possible to show that these alternatives reduce to the same thing, it will be more convenient, at the moment, to think of a reproductive mechanism.

(4) This mechanism may take many different forms and act at many different levels, for characteristically, it evolves. So it would, perhaps, be more accurate to assert a principle of reproduction as our requirement. The point that needs emphasis is that 'Reproduction' is not intended to mean 'Replication', in the sense of creating accurate images of the ancestor. The process of reproduction (at whatever level it is realised) is imperfect and the resulting offspring include variants upon the original.

(5) Thus it is customary to consider an active source of variation as a prerequisite of evolution. In biology, the source of variation acting upon the genetic reproductive mechanism, is genetic mutation.

(6) We remarked that 'Food' or some other conservable commodity must be 'restricted' in the environment. Even this much (the metric over the states of the environment) implies a structured environment. Let us add to this a further structural requirement, namely, that some of the variants shall be at an advantage in the competition for survival.

(7) To elaborate the idea we must be more specific about this 'advantage' and it seems reasonable to introduce a principle that certainly applies throughout biology and which I believe to have universal application. It is that the environment is such that it favours an increased degree of organisation in any process abstracted as a self organising system. This principle implies that there is 'advantage' associated with co-operation. Co-operation involves forming coalitions wherein the participants communicate. In this sense the participants in the coalition form a subsystem or organisation and relations between subsystems add structure to the system as a whole. In biology the process is evidenced by the development of muticellular organisms rendered stable by distance receptors and nervous systems and societies rendered stable by languages and channels of interpersonal communication.

(8) Given these conditions inflow of 'Food' into the 'environment' will initiate a process of evolution, wherein elements, once created, tend to form increasingly organised aggregates and successful variants tend to be selected and reproduced.

2.3. A Specific Evolutionary Model.

The environment is a lattice of the kind shown in DIAGRAM 3. In a more detailed discussion I have considered various topologies (torroidal, infinite plane, and so on) but for the moment we shall restrict the discussion to an infinite plane. Each Node in DIAGRAM 3 is a point at which 'Food' becomes available and where one of the evolving elements or (as it is convenient to call them) automata, can sit and feed. The connections in the lattice are pathways along which automata can move from one node to another.

The structure imposed upon the environment is (i) The set of pathways or nodal connections, and (ii) The rule for delivery of food which though open to rigorous expression, is more conveniently visualised as in DIAGRAM 4., where each node is associated with a Food Bucket filled through a constricted aperture. When an automaton rests at a node it eats food from the bucket faster than this commodity is replaced through the aperture. The electrical current equivalent is also shown in DIAGRAM 4. It is essential to realise that the food supply and the network are distinct and that the food neighbourhoods that may be introduced by making cross connections are commonly not the same as the nodal neighbourhoods that determine relations between the automata.

Rule 1: The simulation proceeds in discrete stages, $t = 1, 2, \ldots\ldots$ at which automata move.

Rule 2: If the amount o ood, u_i, in the ith Bucket exceeds a level an automaton is created at the ith node. This automaton may be of 'type A', or 'Type B' according to whether it is the 1, 3, ... odd th, or the C,2 .... eventh to be created. In a refined version of the model, creation of an automaton at the ith node would be a function of u_i (determining the probability of this occurrence) and a chance variable. The refinement s motivated by analogy with the creation o bubble nucleii in a superheated liquid where the local temperature (by analogy with u_i) determines the probability of a bubble developing, but the appearance of a particular bubble also depends upon nucleii being available. If these are uniformly distributed the model is not greatly changed by this refinement.

Rule 3: Automaton creation has a cost β which is removed from the ith bucket.

Rule 4: Once produced, an automaton, say the nth automaton, must eat food from the node bucket where it rests, at a rate proportional to the amount of food in the bucket. In eating it accumulates a reserve denoted as ϵ_n in an internal bucket. Thus if it rests on the ith node, it eats proportionally to u_i.

Rule 5: Creation of an automaton implies that food has been converted into structural material which has a tendency to decay.

Rule 6: Automata move in discrete jumps, including the possibility of jumping to the same location. The cost of maintaining the structural material of an automaton such as A or B is γ per move and $\alpha > \beta + \gamma$.

Rule 7: If $C > \theta_n$, the nth automaton disappears.

Rule 8: Automata A can move one up or one down and into the same node. Automata B can move one right or one left or into the same node. An automaton can inspect any location into which it can move and determine the food available at this location.

Rule 9: Automata are designed to survive. Consequently their behaviour is determined by a θ maximising decision rule. Consider node a where an automaton rests at t and the pair of nodes to which it may move at $t + \Delta t$, denoted b and c. Let Δu be the minimum difference in local food concentration that an automaton is able to detect. Now, at t, the automaton inspects the food available at a, b, or c and moves at $t + \Delta t$, where there is most food. Many situations are unambiguous. Whenever

$u_b - \Delta u > (u_a \text{ or } u_c)$ the automaton moves to b.

$u_c - \Delta u > (u_a \text{ or } u_b)$ the automaton moves to c.

$u_a - \Delta u > (u_b \text{ or } u_c)$, it remains motionless.

But if $u_b = u_c$ the situation poses an undecidable problem and in this case the automaton must, again, be motionless..

Rule 10: The rate of food inflow is less than γ the rate at which θ must be metabolised to maintain the components of the automaton. Consequently, the automaton cannot survive if it remains indefinitely at one node, and in particular, if it does so because the problem posed by its environment is undecidable.

Rule 11: An undecidable problem is resolved by creating an automaton for which the environment problem may be solvable. If it is not solvable the new automaton will be unable to survive so that on the whole an account of evolution refers to the successful variants.

A new automaton is created by combining original automata. Combination occurs when (i) A pair of automata say x and y, reside at the same node coincidently (ii) at least one of the pair is presented with an undecidable environmental problem. The combination of these automata is written x o y.

The possible combinations of A, and B, are A o A, B o B, A o B and their capabilities are indicated in DIAGRAM 4. The cost of maintaining the fabric of a combined automaton is o. If $o > 2.\gamma$ the cost per unit of material is greater in the combination and if $2\gamma > o$ the reverse is the case. With a few obvious modifications the rule of combination applies at any level of development, for example, to yield structures like A o A o A or A o B o B of DIAGRAM 4 which have elaborate behaviours.

Rule 12: The rate at which an automaton moves depends upon its value of Θ.

Since we are considering discrete stages this rule implies an order in which, as part of the simuluation, the various adjustments are made to the values of Θ and the location of the automata. These adjustments are made starting with the automaton having the highest Θ and taking the whole set of automata in the order of their Θ values.

It may appear that this is a matter of convenience but the model is deceptive in this respect. For, very early in the process of evolution, the outcome depends to a large extent, upon which automata are dealt with before the others.

2.4. Some Results.

In our leisure hours Alex Andrews and I have considered the possibility of programming this model on a computor, but although Alex Andrews has made a draft programme, we have not, as yet, obtained any data. Consequently I shall describe some broad results obtained from a crude paper and pencil simulation. Due to the tedium of this work some mistakes are possible and whilst the important conclusions are based upon results that appear to be repeatable, they should be regarded as tentative until the simulation has been realised on a computor.

As the food inflow builds up a food concentration at the nodes, automata are created, move about, and combine with one another. Whilst the food distribution remains on average, uniform, there is, on average, an advantage in the combination A o B but, in regions that are populated by a particular species of automaton, say A o A, this may not be the case (commonly A o A or B o B would have a better chance of survival than A o B), because the feeding pattern of the prevalent species induces a characteristic pattern upon the food distribution over the nodes in this part of the environment. The mean population size depends initially upon the rate of food inflow (and this remains the case if $\sigma = 2.\gamma$). However, in the arrangement that has been used, $2\gamma > \sigma$ and the mean population size can increase either by increase in the food inflow rate or by the development of combined and more efficient species of automata, with the food inflow rate invariant.

We have already cited the interaction between the automata and their environment, due to the fact that a given behaviour induces a characteristic pattern of food depletion. Now the behaviour of any automaton is a function of its own state and the state of its environment. But

But where this interaction is very strong, due to the concerted activity of many similar automata, the state of the environment is increasingly determined by the behaviour of this population, and as a result, (i) it becomes difficult to speak of any automata in isolation and in practice (ii) the automata form groups wherein the individuals play specialised roles.

At a somewhat earlier stage it seems reasonable to regard the activity of the automata as co-operative, since pairs of automata behave in a fashion that increases their joint chance of survival (the behaviour would be impossible for the individuals alone). The action of the isolated automata that constitute the pair is correlated. The food concentration sensory mechanism has begun to serve a different function (which it reasonably can in these conditions) namely that of a communication mechanism whereby one automaton senses the presence of the other in terms of food depletion. The 'memory' capacity needed to rationalise this statement resides in the intertial characteristics of the environment.

Now, automata, as such, cannot sense a group of automata. But groups can have a characteristic behaviour unlike the behaviour of their components and we, looking on omnisciently, can recognise a group. Suppose, then, that we decide that a group manifesting a certain behaviour shall be at some advantage in evolution. Being omniscient, we are in a position to modify the environment so that this is the case. The simplest procedure is shown in DIAGRAM 5, where the experimenter in the role of nature, determines that the favoured group shall receive a greater food inflow than the rest. He pursues the group about the environment, feeding it. Notice that the procedure of DIAGRAM 5 is a special case of selection due to the constraints in the environment. The actions of the experimenter constitute a sequence of constraints that might have existed in nature to favour this group.

A well fed group becomes increasingly dense. More automata are created within it. Since they are created in a very specialised environment their development, by combination, is largely determined and, on average, is compatible with this environment. Notice that insofar as the group is regarded as an individual entity this process is a mechanism of group reproduction which has evolved. The existing mechanisms of automaton creation and automaton combination are constrained to act in this fashion just as high level selection of lower level selectors is responsible for the amplification discussed by Ashby.(13).

When the density of automata becomes high enough there is a very interesting discontinuity in the structure, reminiscent of crystallisation. Although the group moves about as a whole (modifying the environment where it resides), any individual has a behaviour that is invariant relative to its neighbours. Some individuals, for example, behave as units in a transmission line, composed of a chain of automata, such that motion of one induces food depletion that induces a completely determined motion of the next automaton. Transmission lines have a critical role in the stability of the group since they act like a nervous system that determines a direction of movement for the group as a function of the food available in other parts of the environment (sensed by the relatively unrestricted activity of the outermost automata) and the food available internally (or, more cogently, the internal state).

An impartial observer would be bound to regard the group (not the individual automaton) as an organism, to talk about the behaviour of this group organism, and say that it, rather than its components, made decisions. Although it is occasionally possible to locate a subset of the component automata that remain in a particular critical relation to the others, so that their state determines the group behaviour, this measure of differentiation is uncommon. Ordinarily, there is no place where decisions are made. Rather, Decision Making is a distributed characteristic of the group. In McCulloch's words,(14), the group has a 'redundancy of potential command'. Also, to cite McCulloch again, the components of the group have a 'redundancy of mechanism' for, whilst it is true that differentiation leads to a division of labour, the process is partly reversible. Each component retains the potentiality of acting like an automaton, though it may never exercise this potentiality. Various kinds of disturbance invoke the latent capabilities of the components, and the job of maintaining group stability can be done in many different ways.

The group, being a stable entity, is in a certain sense, aiming to survive. It makes 'Decisions' with this objective. We may thus ask what is maximised by the group 'Decision' Rule by analogy with the Θ maximising function of the Decision Rules of the individual automata. Obviously a prerequisite of any group Decision Rule is maintance of the components. The Θ values of the individual automata must be sufficient. But it is certainly untrue to suppose that mere quantity of food is the most important factor. Indeed for a group the distribution of food (in a part of the environment that the group may visit) is much more significant.

But the distribution (as sensed by the outermost automata) need not necessarily indicate the quantity of food which can ultimately be collected by the component automata. For a group is able to compete with or co-operate with or engulf other groups of automata, possibly having a different organisation. A food distribution indicates the existence of another group. Now our group may be instable because it lacks, say, the species A o A o B and it would be a possible but very esoteric exercise to represent this 'need' for A o A o B automata as a 'need' for a certain distribution of food in the environment. The required automata might be created internally, given a rich enough environment. But equally, they could be obtained from parts of the environment where they existed in any case, or by a symbiotic relation with some other group having a suitable organisation, or by engulfing some other group containing automata of a more elaborate kind, such as (A o A o B) o (A o B). In the latter case, we assume that (A o A o B) o (A o B) automata are unviable in the conditions that prevail within our group. Thus if our group engulfs the group in which they exist, these elaborate creations will come apart, leaving their A o A o B (which we know to be stable in the conditions of our group) in order to satisfy our groups need. Of course, our group may not succeed in this competition. It may be engulfed. At any rate, it seems legitimate to remark that the reasonably interpreted objectives of the evolutionary system change in kind, as well as degree, when the system evolves.

Finally, an observer would experience difficulty in stating the boundaries of whatever he regarded as undergoing evolution. He is at liberty, of course, to doggedly consider individual automata. But if he does so, their behaviour will be increasingly determined by group relations and (unless he admits to this and considers the whole group) his predicitions will become less and less accurate. On the other hand, if he adopts a changeable image, examines sequences of structually different systems which constitute a self organising system, the group of automata he will consider capable of making decisions and manifesting a behaviour not only moves in the environment, but expands. We have tacitly admitted a couple of expansions. First, there is the literal increase in size of the group, increase in density of automata and often increase in volume, as well. Next there is a linguistic expansion.

An observer who talks about 'group decisions' is making statements in a metalanguage relative to the observational language in which he previously discussed 'automaton decisions'. In precisely the same way the message conveyed, at the level of groups, by a food distribution is a metalinguistic statement relative to the message conveyed, at the level of automata, by local food depletion. In other words there is a linguistic expansion that entails the construction of those metalanguages that are needed to maintain the stability of the evolving system. The process can be otherwise described, as the development of relational logics in a domain of abstract objects or as the development of communication modalities in a domain of physical entities with which the abstract objects are identified. The phrasing is a matter of choice. But the choice of a linguistic description is legitimate insofar as that which evolves is credited with the ability to decide.

3.1. Interpretation.

The process can be interpreted as a self organising system at many different levels. Thus one observer may regard individual automata as unitary entities whilst another observer may take groups of automata as unitary entities. In either case the states of the process can be metricised. Ordinarily, the abstract objects in this process will be identified at the same level of discourse with physical objects and the model will be used to make predictions about the real world. As it is formulated, this process describes a mode of evolution which is too general to be conveniently identified with the main course of biological evolution (the formulation is deliberate since we wished to indicate that evolution is not uniquely characteristic of biological organisations). It reflects some peculiar biological products rather well, (the colonial amoebae are a case in point), but it would be necessary to introduce constraints favouring a spatially compact reproductive mechanism in order to generate the mechanism of specialised and sexual reproduction which, as Tyler Bonner points out, is a consequence of this constraint.

3.2. Principles

The evolutionary process gives rise to mechanisms that satisfy the conditions for stability of any organisation given environmental constraints. The dictum that each set of conditions must be satisfied is a <u>principle</u> of evolution in the sense that mechanisms readily interpretable as mediating a common function become apparent regardless of the level at which the process is identified. The constraints that determine <u>how</u> these principles are manifest will be called Evolutionary Rules.

There is, for example, the principle of reproduction. If groups of automata in the model are identified at a cellular level, we indicated how a mechanism of reproduction evolves. If the automata are identified at a cellular level and distances in the network are identified with similarities (rather than spatial displacements), a 'group' becomes a 'tissue' and the same reproductive mechanism mirrors the induction and habituation process whereby the form of a tissue is maintained throughout its development. But this (in Cybernetic terms), is a reproduction of <u>form.</u> Further the comment is useful, since <u>form</u> could be preserved <u>otherwise</u>, by rigid constraints, or by exact replication from a pattern, which would entail the evolution of completely different mechanisms.

3.3. The Paradox of Identification and the Status of the Models

Given a physical object, it can be identified with an indefinite number of systems and conversely, an indefinite number of rational state descriptions correspond to an arbitrary system. This paradoxical situation is commonly but unsatisfactorily resolved by recourse to our normal way of speaking. We can readily construct an identified state description of this table such that the resulting system will 'learn' or be 'concious'. But since we do not normally speak about tables 'learning' or being 'concious' the proposed system is rejected on these grounds.

This dogmatic approach is essential. Nevertheless it is dangerous. For what one generation of scientists regard as unable to learn is often viewed as able to learn by the next and precisely the same comment applies to distinctions such as able or not able to reproduce, alive or inanimate, intelligent or devoid of intellect, which rest at the foundation of each evolutionary principle. The conventions of identification which are built up may seriously hinder the advance of science although they are needed for the conduct of everyday experimental work. Nowadays, for example, we are bedevilled by conventions that lead us to look for such things as 'memory', 'personality' and decision making in a particular place or a particular process and the evidence of models is needed to demonstrate that although we can occasionally locate the vaults of memory or the seat of decision making, this discovery is exceptional.

The point about a model (one of the evolutionary models of the kind we have considered) is that state descriptions of it are restricted (and consequently possible identifications between it and a self organising system are restricted). Initially there is always a unique state description. In the case we have considered it might be the food distribution stated at t = o for

'A Plane Lattice with M nodes'

or

'The Lattice on an infinite plane'.

If 'M' is large or the plane infinite, it would soon become impracticable or impossible to use a unique, omniscient, state description. But certain rules still apply, for example, definitions of the forms of automaton, and in principle, it is always possible to work back through the sequence of transformations and reach the state description which is unique.

Further, because of this we can comment upon the status of a model with reference to a principle it exhibits. Consider the principle entailed by the assertion that some organisations, especially man, control the environment in which they develop. The present model is able to demonstrate the discontinuity in evolution that occurs when this statement becomes true.

But the status of the model, with reference to this principle, is vastly improved if, as in DIAGRAM 6, we introduce the possibility of effecting permanent changes in the rules that govern the 'environment, in DIAGRAM 6, by altering the width of the local food inlets, depending upon how much has been eaten in this locality.

The change induced by the behaviour of one generation of automata is neither reversed by disuse nor by the action of the next generation.

3.4. Review of the Evolutionary Principles we have Considered.

The inbuilt characteristics of variation, reproduction, competition and co-operation, become mediated by mechanisms which arise as the products of the evolutionary process, and to a large extent, supercede the original mechanisms. But the originals are potentially available and a single mechanism rarely serves a single function. Thus there is redundancy of mechanism. The distinction between an organism and its environment is obscured due to the expansion of the organism when it gains control of its surroundings and because of the concomitant increase in coupling between the organism and the environment, we cannot say, unambiguously, what evolves. Further, it is legitimate to assert (as Tyler Bonner does, in the case of biological evolution)(15) that whenever a species evolves so also does a miniature process of evolution (or sequence of interactions between the organism and constraints induced by the organism in its immediate environment) which we call the development of a member of this species. On each score, indeterminacy of spatial and indeterminacy of temporal extension, there is a redundancy of potential command. Viewed as data processing or computing arrangements the different evolutionary forms, or their replicas in the miniature sequence of ontogeny are characterised by an hierarchically ordered sequence of languages. Further the different orders of organism coexist for evolution is a provident business that uses outdated parts to build up to-date forms, rather than starting afresh with each variant. Stability requires that the co-existent forms can interact or, in terms of data processing, that the different orders of language characterising these different forms, are translatable.

Discontinuities, engendering differences of kind rather than degree are apparent to any observer who adopts a given reference frame. Of these we have mentioned: (1) The point at which the group organism, if consistently reinforced, increases in density and acquires an internal communication structure. (2) The point at which the evolving organism can exert sufficient control over its surroundings to largely determine the environmental constraints that catalyse the differentiation of the species. At this stage, the evolution of the species becomes autocatalytic. (3) The point at which the objective for which individuals of a species are competing changes from 'food' to 'members' of a different species. (4) The point at which organisms of a species, already able to register their activity by a permanet impression upon their environment, become capable of conveying experience from one generation to another generation. This is one of the most important events in the process of evolution, for differentiation and adaptation by selection is replaced by a far more efficient method.

4.1. Practical Utility

Has this kind of abstraction, or indeed, has any theory of a self organising system, a practical use? The reply is in the affirmative. It was given, most comprehensively, when Stafford Beer (16) spoke from this platform at the last meeting of this Association. You will recall that he conceived industrial organisation as biological. Any other form of organisation would be or, in view of technical progress would become, crassly instable. Speaking on the same theme, Georges Boulanger(17), our President, envisaged industrial controllers, like brains rather than automata, able to plan what job to do (not merely how a prescribed job should be done).

The crux of the matter is that man has reached a critical stage in his own evolution. Technical advances offer him unlimited power to control his environment. In order to stabilise a system which (because of this) is subject to rapid changes in state, men and their social aggregates must be closely coupled. But the price we have rightly and necessarily paid for an advanced technology is division of labour and this, on the face of it, seems to prohibit close coupling. For the existence of the physical means of communication is not enough to ensure that communication will take place. The real limit lies in lack of compatibility on the one hand between men themselves (when it is a question of mutual understanding, and being intelligible) and on the other hand, between man and the automata (or automaton-like organisations) that form an increasingly dominant feature of an increasingly man-made environment.

Specialised men speak different languages. They do not share a universe of discourse. Specialised machines are only slightly more distant, in a semantic sense, than some of our own species. One often hears that there is a need for wisdom and for breadth of mind, and if these words augur or describe the ability to translate between diverse tongues and the flexibility to attempt different kinds of translation, then we must surely agree with this hoary sentiment. But it is a Cyberneticians job to inject the wisdom, to make certain that translation and understanding actually occur. As Beer pointed out, the ideal organisation is biological and since we are thinking of man it is man-like.

Take man as a model. Much of human behaviour is determined at a reflex level, by biological automata, some of it is thalamic, some at the other extreme, is cortical. But there is no discontinuity. Although the field of consciousness (which I take to mean, at least, the universe of discourse we have in common with our immediate neighbours) normally refers to a facet of cortical activity alone, it can be otherwise (when physically embarassed, for example, or in stressful conditions). On these occasions we do consider our breathing, we do become aware of our heart rate. Obviously translation is possible when it is necessary, between the normally disjoint languages that characterise different levels in our nervous system. Now the ideal environment would preserve this much continuity. It would extend our internal organisation into a single self organising system, which, in a sense, must be an image of ourselves.

The problem of constructing this kind of environment is enormous. Except in some special cases, building it is out of the question. The best we can do is encourage the system to evolve.

4.2. Particular Applications

The theory is also applicable to detailed projects. In our laboratory we have been concerned with the interaction between men, participating in Decision Making groups. Our hunch, which is discussed in a separate paper, is that the kind of optimum decision making that is credited to the joint action of a group and is deemed superior to a mere concensus of individual choice, is a product of redundancy of potential command that can only occur if the group is a self organising system. To realise such a system the communication structure, which in most experiments is invariant, must be made labile and continually changed as a function of the decision making activity of the several participants.

We are also concerned with the interaction between man and various adaptive automata used for teaching skills and for aptitude testing. In this case we are anxious to distinguish between the common form of man machine interaction (where the characteristics of the machine are unaffected by the man's behaviour) and the present arrangement in which these characteristics are adaptively modified.(18) Given this modification, man and automaton become closely coupled and the logic of the joint system (which is a self organising system) is isomorphic with the logic of a conversation in which concepts evolve. This work, also, is described in a separate paper.

Finally, we are interested in cognitive automata able to learn about their environment. Since this project is at a very early stage, there are few tangible results and I shall use some preliminary thoughts on the subject to illustrate the way in which ideas of evolution or of self organisation come into a research programme.

4.3. Evolution of Cognitive Machines.

Current thinking about Cognitive Machines is divided over the issue of specification. On the one hand, it is possible to construct automata that imitate the perceptual filters of various animals. For the visual modality J.Z. Young has described Octopus Vision whilst Lettvin and Matturana have displayed the attribute filters of the frog and revealed a mapping of fibres from the retina to the colliculus that bears out the earlier predictions, made in terms of a network of artificial neurones, by McCulloch and Pitts(19). Nor is it necessary to imitate a particular animal, for the mathematics of perceptual filters has been worked out with nicety chiefly by von Foerster's department at the University of Illinois(20). Although the frog does not appear to learn any percepts (and its filters are invariant) this is not a necessary feature of the approach (certain parameters of an attribute filter can be changed by adaptation without losing the point of the filter model). A very well known but special case is Uttley's (21) conditional Probability Machine where the categorising network is invariant but the contents of state probability registers or, equivalently, the impedance of the network connections are variable.

On the other hand, there is a school of thought that advocates a completely unconstrained network in which some kind of perceptual filter is built up by selective reinforcement of signal pathways or connections in a network. The reinforcement is administered by an arbiter who views the moment-to-moment performance of a device such as Rosenblatts Perceptron(22). One difficulty is that no realistic machine can be unconstrained. The perceptron has initial and possible connectivities determined in an arbitrary fashion 'by chance'. But these, in any particular case, specify a definite structure. It is argued that over many cases the form of structure will be irrelevant to certain statistical parameters of the networks behaviour.

Alexi Ivanhenko (22) has examined the problem from the point of view of his combined system theory and it seems as though his work will do much to reconcile the structured and the 'unstructured' schools of thought.

The far from original point of view we have adopted has no connection with servo theory. It can be interpreted as a compromise between the extremes or alternatively, as a completely different approach.

4.4. An Approach Involving Maturation

Cognitive Machines, able to learn as much as you please, are highly constrained yet (barring omniscience) they are unconstructable. Let the automaton be a system evolving in a network. If we call the whole sequence of development a 'Cognitive Machine' its behaviour is virtually determinate and constitutes a self organising system (when viewed by an observer in a suitable sequence of state descriptions). But, as we noticed previously this observer does not adopt the attitude of a constructor. He does not determine the state description as a constructor would determine his. Instead, he reaches a compromise in order to make sense of the behaviour. If, on the other hand, we call the final product a Cognitive Machine it has a virtually chaotic behaviour and we are in the position of J.Z. Young with Octopus or Lettvin and Matturana with their frog. We must look at the behaviour of the final product and discern a state description to fit it. In the case of the frog, a fitting state description turned out to be a bizarre set of attributes, ideally suited to a world of flies and shadows (indicating predators). Our cognitive automata will have a very different world, no doubt, and they will not always be as tractable and structured as the frog. But we shall no more make sense of even their simplest behaviour by classical analysis (our preference) than the earlier workers made sense of frog vision by shining discrete points of light on the frog retina.

The mechanical basis of the evolution of a Cognitive Automaton is a region of activity in a network of adaptably connected aritifical neurones. Such networks have been described and simulated by (24) R.L. Beurle and, in a simplified version by Farley and Clark (25). Whilst R.L. Beurle was chiefly concerned with a model for cortical activity and determined the initial constraints from physiological and anatomical data (he made the artificial neurones with characteristics similar to real neurones and specified the possible connections according to the statistical connectivity of a real cortex), we may take a broader view, unrestricted by the parameters of any actual brain. But, in other respects our network will resemble Beurle's. Thus, like Beurle's network, there will be a restriction upon the number of active 'neurones' per unit area (consequently, action will be competitive). Further, the activity of any one neurone will depend upon the correlated activity of several others, inducing co-operation. The changes in 'synaptic' impedance which determine the effect of activity at a distance upon the state of a given neurone will depend upon the history of coincident activity and the previous values of a reinforcing variable.

Let the individual neurones become indefinitely small. The network becomes a transmission medium in which waves of activity are propagated. Adaptation at the synapses alters the impedance of this medium to different forms of wave. Coincidence of waves may, in certain conditions, give rise to a novel wave unrelated to the form of its ancestors, but normally the waves are replicated. Finally, we require that the network is autonomously active, and that in the absence of external constraints (or inputs) the variety of the oscillation within it is maximised. Inputs from the external environment induce waves of activity as suggested in DIAGRAM 7, and (in some conditions, when there is a correspondence between the input wave and the existing oscillatory mode) these will dominate the autonomous activity. Outputs are obtained, as in DIAGRAM 7, by transforming a wave front into a signal, using a special component or a wave attenuating region built up from a block of the artificial neurones.

Recall that there is an isomorphism between the evolutionary model we discussed previously and this kind of artifact. An automaton or group of automata acting like an organism corresponds to a 'mode of activity', or, loosely speaking, an 'active region' in the network. Actually, several correspondences are possible, for example, we can obtain isorphism by mapping the automata in their action metric onto artificial neurones in a metric which is a function of signal proximity and internal state. But for the present purpose these alternatives are less convenient and we shall assume that active regions are the entities which compete and co-operate.

This correspondence was first pointed out by J.W.S. Pringle (26) who aimed to show that learning in the mammalian brain was a kind of evolution. His paper, demonstrates an identification between cycles of typifying activity and behavioural invariants characterising types or species of organism, and should be consulted. For the moment, it will be sufficient to remark that the features of evolution we noticed in the case of our automata make their appearance also in the case of active regions (some recent work by Ashby indicates that this is no quirk due to our choice of the parameters but that the evolutionary process is characteristic of any large dynamic system with many metastable states). However, some care is needed when interpreting the phrase 'active region' (in particular when talking about reproduction). For an 'active region' is neither activity at a given location in the network alone nor the plastic changes which are a consequence of this activity alone. It is due to the interaction of these only partially seperable entities.

The active region moves around in a network but its activity, or certain facets of its activity, remain invariant. Thus it can be recognised or, conversely, since a pattern of events remains invariant although the underlying activity is differently embodied, this pattern is reproduced.

Thus there is a satisfactory analogy between a biological species and a pattern, and between a biological organisation and one active region. When we say reproduction occurs, we mean that a form of behaviour, characterising a species, is replicated. The process of replication which is the ontogeny of the organism, is characteristic of an organism of this species, and we have seen how the evolutionary process gives rise to this mode of development. Similarly, the history of the species, its phylogeny, is a comparable story told at a different level of discourse.

The evolutionary system is an abstraction of the process identified with states of the network in such a way that relevant events remain distinct and insofar as a coherent behaviour of the system can only be obtained by telling stories at different levels of discourse, it is a self organising system because an observer who tells a coherent story must adopt a sequence of different state descriptions. That it is a self organising system rather than a haphazard system is guaranteed by the condition that the stories are comparable.

D.M. MacKay,(27) who first considered this field, remarked upon the importance of hierarchical structure in any cognitive network. His artifact was a device of the kind we have described. Suppose, for simplicity, that it receives no input apart from the reinforcement variable though the argument is unaffected if the input dimensionality is increased. In this case, the artefact agitates its environment by making trial actions. It is tempting to conceive the device as dissatisfied with a changeless environment. The autonomy condition ensures that it must attend to something. Failing any change in the reinforcement variable, it tries out fresh modes of internal activity which induce fresh sequences of trials.

Assuming that some of its trial actions are reinforced there will be a tendency for structuring to occur. The active regions that select the reinforced trial action becomes stabilised and reproduce. Thus, it is legitimate to consider stable sequences of actions which can be selected as a whole. The active regions that induce the stable sequences of actions interact with other active regions that are incapable of selecting actions as such but select, instead, from the set of action sequences. Consequently, some of these are reinforced and become stabilised and reproduced. As MacKay points out, a process of this kind yields a selective hierarchy in which the organisations at different levels, selecting trial actions and sequences of actions, and so on, are characterised by distinct languages. Thus, to talk about (or to make selections in terms of) an action sequence, is to make utterances (or selections) in a metalanguage with reference to utterances or selective procedures that concern the individual actions. We have cited this characteristic, also, in the case of the evolution of automata.

The hierarchy is an informational concept and it can assume many different physical realisations. Invoking the principle of redundancy of mechanism we might predict that hierarchical structure would be multiply represented. One realisation, for example, is a literal structure of levels, like an adaptive controller, in which a high order controller determines the parameters of a lower order controller. Another realisation is a separation of levels by phase displacement of messages (the kind of organisation proposed by Crane for his Neuristor Networks). (28) Again the separation may appear at the level of components which are selectively sensitive to different aspects of the same physical signal (one component might respond to the average frequency of impulses whilst another, developed by evolutionary modification, might respond to phase difference or interval between impulses).

In any case there is no doubt that any machine which purports to learn about a structured world (or to recognise aspects of its structure) must possess an informational hierarchy (if it did not, the adaptive process would be too slow to be realistic). Further the structure of the hierarchy must be related to the structure of this world and since building an hierarchy is tantamount to the process that occurs when an unstructured machine adapts to the world, it is argued that the hierarchy, or its main outline at least, must be built into any realistic machine, either mechanically, by a predetermined connectivity between the elements, or informationally, by a heuristic programme.

Newell, Shaw and Simon, and Marvin Minsky have pursued the heuristic approach, Oliver Selfridge is more concerned with the mechanical approach and mathematically they are largely equivalent. Minsky and Selfridge have indicated the basic mathematical requirements in a recent paper.(29)

One practical attempt to meet these requirements is Oliver Selfridge's recognition programme 'Pandemonium'.(30) The programme determines a set of computational routines called 'Demons' that recognise specific attributes of their environment. The lower level 'Demons' in an hierarchy send recognition signals to demons higher in the hierarchy that form 'percepts'. In turn, data regarding 'percepts' is collected and a state of the environment is defined. Each component is selectively reinforced according to the agreement between the designated state and whatever designation an omniscient arbiter, also aware of the environment, has in mind. Depending upon their value in obtaining reinforcement, demons at a lower level are preserved or rejected by higher order demons. If a 'demon' is rejected at least some of its component subroutines are used in constructing a further routine or 'demon' to replace it. Since the process is selective and provident of variation, it is evolutionary.

Now a device like 'Pandemonium' is certainly 'structured' and is capable of development within this structure. But it seems to me that this is not the only way of introducing structure and it is not necessarily the best way. It is particularly efficient if we aim to evolve something that imitates a certain transformation of input data, like a passive network that selects one output for each combination of inputs. Now a frog brain, say, is a machine that resembles a passive network, although it probably evolved from some brain that was not. But a frog brain does very little learning whilst the cognitive Machines that are of the most interest, learn a great deal. Thus, without denying the propriety of this approach, the word 'Cognitive' seems to be used minimally or even trivially in connection with pattern recognisers, for if we take the evolutionary viewpoint seriously (as it is indicated, for example, in the paper by J.W.S. Pringle) a Cognitive System functions because it evolves. It would be absurd to consider it apart from its evolution or (because of the continual interaction between the system and its environment) to separate it from the maturation and adaptation of the brain or other network in which it evolves.

Since we aim to achieve this broader connotation of the word, we shall adopt a definition consonant with the initial contention of 4.4. A cognitive Machine is the environment, usually an internal environment such as a network or a brain, wherein Cognitive Systems are induced to evolve. Design of a cognitive machine thus entails specification of the internal environment together with Rules of Evolution that constrain the development of active regions.

The definition is loaded with philosophical overtones. For the moment, let us put these aside and avoid any detailed consideration of the function of a Cognitive System in order to concentrate upon the practically important issue of determining the Evolutionary Rules.

(1) The 'Cognitive System' is a self organising system. To make sense of its behaviour we must adopt a sequence of different state descriptions. Equally, seeking to constrain its evolution, we must use constraining operations that are pertinent within the current state description.

There is an immediate departure from the existing structural and heuristic techniques since these constraining operations are commonly valid only within one state description. Yet the more flexible procedure is familiar enough. It is well known that different kinds of influence have different potencies according to when, in evolution or development, they occur. Sometimes the precise moment of a constraining operation is critical (imprinting of data in birds) and commonly the pertinent form of a constraining operation is changeful (the entirely different training techniques we use at varying stages in the maturation of a child's brain). At a more intimate level, chemical, mechanical, and neural influences are differentially effective constraining operations.

(2) The Cognitive System regulates its own development and our interference must, so far as possible, be consonant with this regulation.

Thus, the Cognitive System, evolving in a brain, determines the behaviour of an organism. As the system develops, the organism becomes capable of actions that elicit increasingly elaborate situations in the external environment. Now these situations determine constraints, and the consraints induced at any stage are cogent. Actions which elicit other constraints are not embodied in the repertoire of the organism.

We do not pretend to know what constraining procedures should be adopted in the case of an arbitrary Cognitive Machine, however suggestive analogies with the brain may be. But this is where evolutionary models, like the population of automata, become useful.

For it is not too difficult to visualise procedures that alter the course of evolution of a species of automata. These procedures, altering the structure, the food available and the basic evolutionary rules for combining automata, can be transformed into often less obvious equivalent operations upon a network.

5.1. Speculative and Philosophical Comments

In conclusion, let us indulge in some moderately speculative comments about cognition itself, about consciousness and its relation to cognition, about the location of consciousness, and finally, about the extent of and the location of memory.

5.2 The Character of Cognition.

My hunch is that cognition, in its non trivial sense, is associated with a special kind of redundancy of potential command. In the situations that McCulloch envisaged when he coined this phrase, there are a number of decisive elements so coupled that resonance occurs and, because of this, no one element can be credited with making a decision. Now, in the cognitive machines we have considered there is an hierarchical structure, as in DIAGRAM 8, of a kind that Jack Good, with charming and illuminating irreverence, called 'Onion-Like'. The layers of this onion could, for example, be levels of 'demon' in a Pandemonium or they could be any structures, characterised as in 3.4. by different metalanguages. It happens that in Pandemonium a 'master demon' can be credited with decisive action in the cognitive task of categorising an image, providing we 'freeze' the system at one stage in its evolution. But were the stages in the process of evolution less orderly and 'freezing' impossible, (which would be the case if the 'Demons' become active regions evolving in a partly independent fashion) this would no longer be true (for we could no longer specify the form of evidence provided for the 'Master Demon'). Briefly, in an 'Onion-Like' structure, Redundancy of Potential Command is likely to occur, (1). Because of interaction between the decisive elements within a layer of the 'Onion', thus introducing uncertainty regarding the form of evidence provided for the higher layers. (2). Due to communication <u>between</u> the layers which can only occur if the characterising metalanguages of the layers are translatable but which, as we pointed out in 3.4. <u>must</u> occur if the onion represents a stable evolutionary system. There is an interesting feature of the latter kind of interaction which we shall return to in a moment.

Now the special form of Redundancy of Potential Command which I believe to be correleated with Cognition occurs when the decisive elements McCulloch envisaged, namely the components of the resonant system, are <u>ancestors</u> of the resonant system. A moment's consideration of the principles advanced in 3.4. will convince you that this situation is likely, if not certain, to arise whenever the system is evolutionary. But in 4.4., on grounds advances in 4.3., we defined a Cognitive System as necessarily evolutionary.

This is my hypothesis. Whenever a brain deals with the problems imaged by its sensory receptors we can say, equivalently, that active regions evolve in the <u>internal</u> environment of the network called 'a brain'. As a result of this evolutionary process, the internal environment is modified as we noticed it must be in the case of the evolution of automata, so that it becomes adapted to the requirements of the dominant species. The evolution that occurs at this level constitutes the mechanical background of 'learning' in the species of organism that embodies this network as its brain.

There is a basic interaction between this and the evolution of the species of organism in the external environment for certain variants will have brains that favour the internal evolution of intellect, that are "able to learn". Insofar as the behaviour, which is a product of this intellect, permits language and communication, these species have a selective advantage and these brains will be preserved, and reproduced as part of the organism. But in a sense we have discussed the selected brains embody their evolutionary history and recapitulate it by their maturation. The layers of DIAGRAM 8 correspond to internal environments that favour the more primitive species of active region that characterised the intellect of the more primitive ancestor organisms. Similarly the active regions that arise by internal evolution will be ancestors of one another. As development proceeds one or another species will dominate the intellect. But I believe that Cognition occurs as the product of interaction between the species that inhabit these essentially historical layers.

As Bishop (31) pointed out, the anatomical data supports this view. The mammalian brain is an hierarchical structure in which the anatomically defined levels correspond to more primitive brains. In some cases functions appear to have been lost (the thalamus has lost its motor outlet) and in other cases there has been specialisation (the cerebellum). But the mammalian innovation of a large cortex is more accurately described as being essential to than as being above these other brains.

Finally there is an interesting correspondence between "attention" and "cognition", for the "attention" of a cognitive system has a double meaning. On one hand, the organism embodying the Cognitive System, 'must attend to something' and we recognise that the system has a lower as well as an upper limit to its decision capacity. The mechanism that underlies this aspect of attention, in our networks, is whatever mediates the autonomy condition of 4.4. On the other hand, there is a sense in which 'attention' means 'awareness' of a cognitive task and the coupling of subsystems to form a system which, in our view, produces the behaviour that we expect from this species of organism, (if a man is asleep we do not say he is functionless but we do say that he is unaware of his surroundings, for although he may respond his behaviour is not, in our view, characteristic of a man). The mechanism underlying this aspect of attention mediates interaction between the layers (and in man can possibly be correlated with the reticular formation and its connections).

5.3. Cognition and Conscious Activity.

As MacKay has often argued,(32) the statements we make about consciousness have no direct relation to those we make about the mechanism of artifacts. If you decide to say a machine is conscious this decision is a somewhat personal matter, for your attitude towards the machine, which will be completely different if you regard it as a conscious machine, does not arise as the result of a logical conclusion you have arrived at upon the evidence from any well defined set of experiments. Broadly speaking, you require evidence of similarity between you and the machine and whilst the tests you make are perfectly explicit, your criterion of similarity is personal.(6) My view is that we should say of 'X', that 'X is a Cognitive System' as a prerequisite of saying 'X is conscious'.*

This does not imply that any Cognitive System is conscious (indeed, as above, statements about Systems and statements about consciousness are distinct), nor does it delimit a category of conscious objects (you may call this chair 'conscious' insofar as you can construct a set of state descriptions in which the chair is a Cognitive System). On the other hand, it does allow us to restrict the operational statements we make about consciousness to those which could reasonably be made about Cognitive Systems. So, for example, it is unreasonable to talk about the location of consciousness because, operationally, it is impossible to locate a Cognitive System.

My view in this matter stems from the fact that a Cognitive System is able to make self referential and interpretative statements, and is the least organisation with this capability.

The phrase, 'Self referential statement' has its usual connotation. Burke,(33) considering the evolution of Von Neuman automata, argues that the system as a whole must make self referential statements, which interpret the activity of its constituent subsystems, in order to evolve. His requirement is equivalent to the condition of 3.4 that the metalanguages, characterising different layers in the hierarchy of a stable evolutionary system, must be translatable. 'Interpretative statements' are comments upon the relation between the Cognitive System and objects in its universe of discourse.**

* Cognitive Systems that are not similar to us must be commoner than those that are similar. They may be more or less stupid than we are. Those that are not less stupid have a promising application in controlling bizarre environments in which we should be unable to detect any regularity.

** With reference to our previous comments about attention, a Cognitive System is the least system able to construct a universe of discourse.

Returning to the point that our decision to call 'X' conscious, rests upon evidence of similarity, it seems to me that whatever criterion of similarity is adopted, the tests of its satisfaction entail 'X Referential Statements' or 'interpretative statements' (On the part of 'X') about objects in a universe of discourse that is common to X and the tester, (the tester will need assurance that X appreciates these objects in much the same way that he does). Consequently, in order that a criterion of similarity be tested, X must be a Cognitive System.

Further, when a Cognitive System evolves (implying that the physical assembly with which it is identified increases in size or structure) the possible 'profundity' of interpretative statement must also increase (this is only another way of indicating that higher order metalanguages must be constructed to achieve stability). The more evolved the system, the more wisdom is needed to keep it viable.

The word 'profundity' is precise if we restrict our attention to the special case of interpretative statements used to specify a system. In order that we should count a collection of statements as 'the specification of a system' we must, I believe, credit the source of these statements with consciousness. So we are considering the special case of a Cognitive System, agreed to be conscious, specifying a system which is a mutually intelligible image of some assembly open to join inspection. In particular the specified system may be identified with that collection of physical parts that we have identified with the cognitive system which is producing the specification, and, if so, this cognitive system is uttering self referential statements.

Now, in these conditions, the greatest possible profundity of interpretative statement is measured by the degree of organisation - a redundancy - associated with the most elaborately structured system that can be specified.

5.4 The Spatial Extension of the Evolving Cognitive System.

Since there is no absolute dictum about coupling between a pair of systems, there are no absolute boundaries which define the spatial extension of a cognitive system.

Where, for example, does a man have his limits. Even if we have agreed to speak informationally (rather than speaking in terms of anatomy or of energy) these limits are tenuous and changeful. Is his brain a closely coupled system 'the man' as separate from his receptors or effector mechanisms? Or should we consider his influence upon organisations in the environment which may, of course, include 'other men'? We need not dwell upon the issue (since it was discussed in 3.4.) except to comment that the limits depend upon us and upon our choice of system and our objective in constructing it. But, however arbitrary, these limits do exist.

Now, from 5.3., we can also ask, 'What are the limits of the system that a cognitive system specifies as an image of itself'? In particular, 'what are these limits if the system is stable and is evolving'? In this case, the reply must be that the Cognitive System has an unlimited spatial extension. For the more extensive it is the wiser it must be, and the wiser it is, the more states of the environment will be involved, and, commonly, the greater its spatial extension (the greater the region of potential control). You may or may not find this circularity illuminating. It means that the environment needed to sustain the stable evolution of a self organising system (in particular a cognitive system) is, as we argued in 4.2. a similar self organising system.

The case of an unstable system is trivial, because it is unobservable. The case in which the system does not evolve is the special Cognitive System called a 'Perceptual Filter' or 'Recognition Device'. Its spatial extension depends upon the form of receptor whilst the set of states of the environment, involved in the system, are those required to specify the percepts of the system.

5.5. The Temporal Extension of the Evolving Cognitive System.

Let us take it that there is no problem of memory capacity, for any brain like assembly is made up of components such as neurones that can and must exhibit many kinds of hysterisis. True, there is a tendancy to single out one sort of state description, for example, to regard neurones as almost binary registers, and to compute a memory capacity in terms of the limited forms of hysterisis that the resulting system can exhibit. But manifestly this is no more than a convenience and a convention. Neurones have chemically distinct states and mechanically distinct states. The membrane in the region of a synapse is highly structured and could determine memory at a mollecular level. Maybe other bodies than neurones are significantly involved in memory.(24)

The point is, one cannot do anything to a brain, stimulate it or change its chemical environment, without inducing persistent changes of state. This would be true of any large system, with many equilibria, as Ashby points out. But the brain is more like a set of these large systems, each defined in 'Binary Impulse' terms, or in 'chemical state' terms or in terms of active cycles, and these systems are closely coupled. So that if, for example, neurones are set into cyclic activity (which amounts to 'memory' of the stimulating event) plastic changes occur at the synapses and mollecular changes occur at the membranes and these also constitute 'memory' of the event).

The problem of memory is thus a Coding Problem, for information is registered and processed in many different, closely coupled, and coincidentally utilised representations. Brains are physically capable of retaining an index of any occurrence. But the indexed data may or may not be decodable.

I wish to make only a single point in this connection, namely, that the temporal extension of cognitive systems evolving in hierarchically structured brain like networks (with components as messy and versatile as we suggested) is unlimited. To illustrate the point consider an active region at Layer 1. (higher, in a hierarchy of metalanguages, or historically more advanced than Layer 2), so that a Layer 1. active region can interpret the active regions of Layer 2, in the sense of constructing a system to represent their states. Notice that constructing a system is precisely the decoding procedure needed to extract some arbitrary memorised event.

Now we may interpret the active regions in Layer 1 and in Layer 2, at, say, the level of impulse transmission and abstract them in this sense, as little Cognitive Systems. We admit, of course, that plastic changes also occur, but regard these as irrelevant to the activity. (If they had seemed relevant, we could have examined them in detail). So in this picture, a Layer 2 Cognitive System computes data and hands it on to Layer 1 and some memory resides in each layer.

Now a Layer 1. Cognitive System, defined in terms of impulses, can only make statements about impulses in Layer 1. and the same comment applies to Layer 2. Cognitive Systems in Layer 2; But a Layer 1. Cognitive System can build its own system to represent Layer 2. activity. It is free to select a different representation (say, in terms of plastic changes at the synapses) to the one we happened to choose. Indeed, higher level systems can interpret lower level activity in a variety of different ways (each of which amounts to a coding procedure). Since the process is cumulative and since the evolving cognitive system can construct hierarchical levels it can, in principle, decode from any state of the brain. But there are indefinitely many state descriptions and thus of states. So the temporal extension of the system is unlimited.

Equally, memory may involve the environment. Given a minimal coupling to a minimally adaptive kind of environment it is still true that any change of state in a Cognitive System exerts some effect upon the state of the environment and it is always possible that a system capable of appreciating the significance of this change will evolve. In more adaptive surroundings bits of the environment may be used like bits of a brain. We have seen this occur in our abstract model and there is no sparcity of cases in biology. The bird builds a nest. The nest provides a feedback stimulus that evokes the behaviour of egg laying and the hormonal changes needed to lay eggs. The eggs provide the feedback stimulus for sitting. Because the bird sits the eggs will hatch which provide a further stimulus. So functionally the bird does engulf its environment. But if the organism that embodies our Cognitive System is part of a larger self organising system, which a stable system must be, its memory is distributed and unlimited and resides in language and in systems of evolving concepts.

It seems to me that a human being must be described as a Cognitive System. With Heinz von Foerster I insist, as a matter of faith if you like, that there is a relation between any human being and any other and anything he would call his environment, such that the systems he may build are structured. For, if so, there is something to perceive. Finally, again as a matter of faith, I have the healthy materialism to believe that systems are identified with a physical world which, unlike a metaphysical or merely intellectual world, is open to an indefinite number of interpretations.

These seemingly innocent requirements lead a Cybernetician to some odd conclusions about the immense permanence and immutability of form in such a self organising system.

(1) With Huxley and others I must conclude that 'Me', my 'Consciousness' is not localised. It cannot, for example, be found in my head. Ideally, if not necessarily, the vehicle of this consciousness, a Cognitive System, has an unlimited spatial and temporal extension. It is doubtful if one can say 'my' Cognitive System. There is an intruiging sense in which 'my' Consciousness may develop because the system ramifies.

Because memory is hewn in the stuff of a physical world identified with indefinitely many systems, nothing thought or felt or done can ever be forgotten by the system in which we live and of which we are part. That it should not evolve is inconceivable.

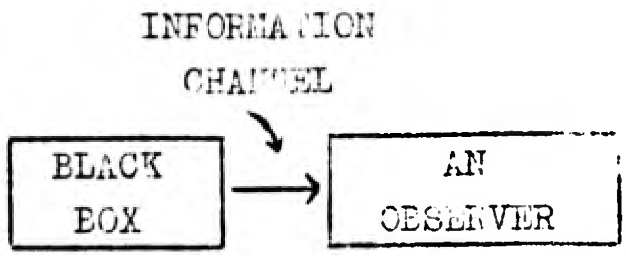

DIAGRAM 1.

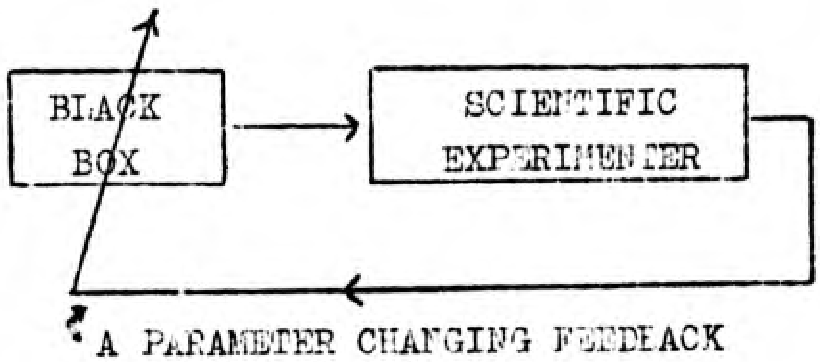

DIAGRAM 2.

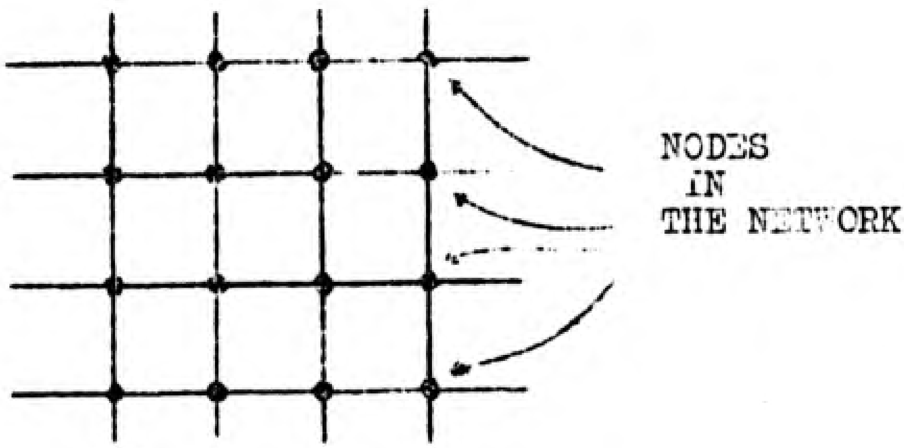

DIAGRAM 3.

(i)

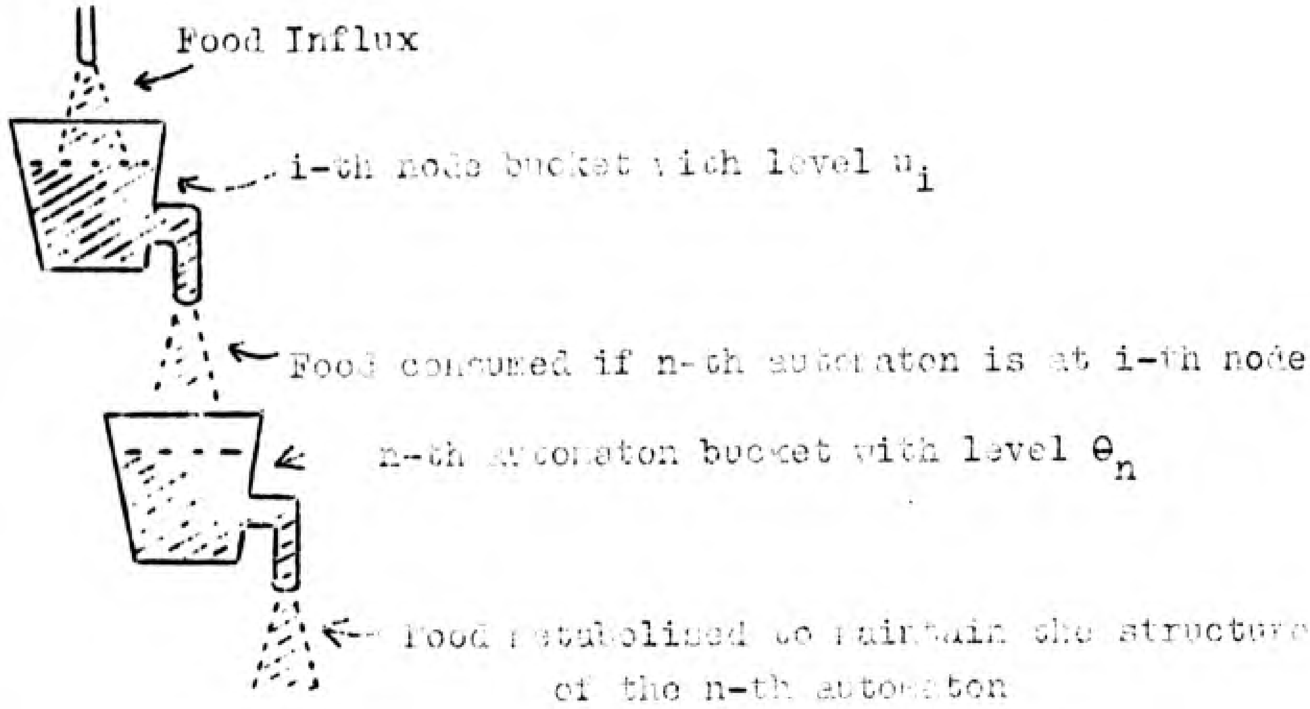

(ii) Electrical equivalent:

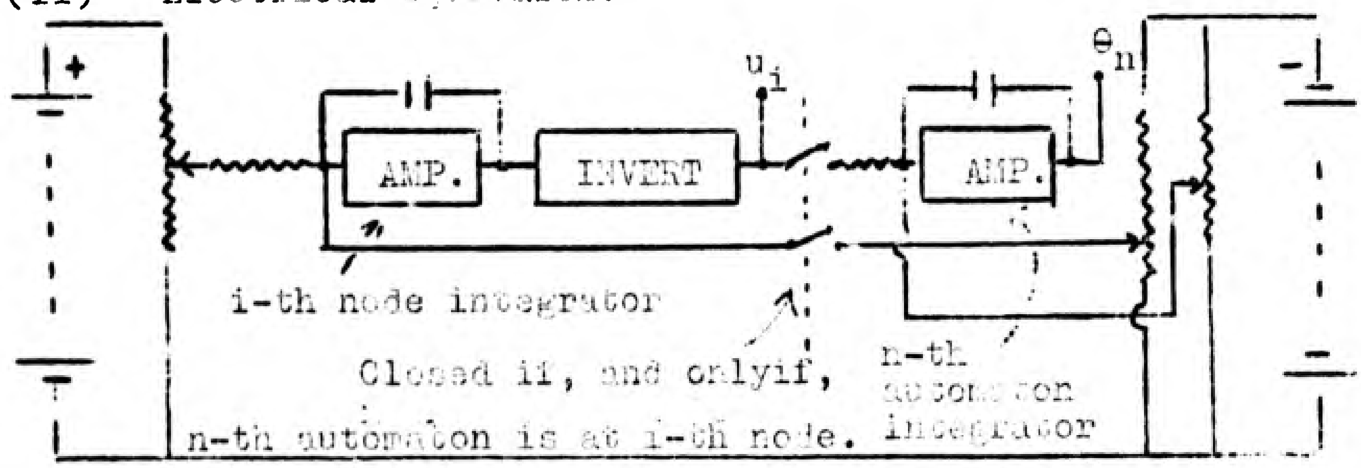

(iii)

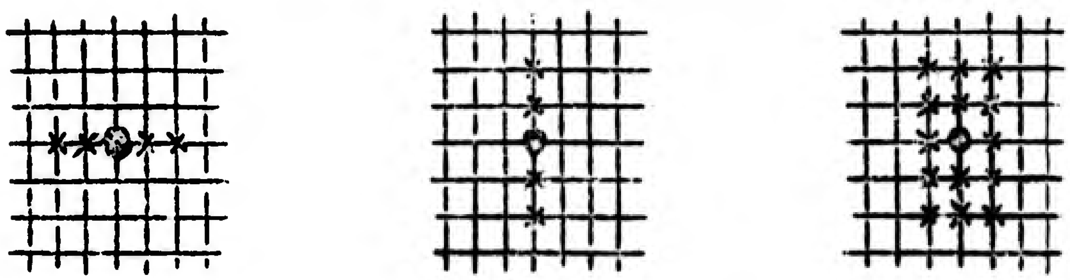

A o A B o B A o B o B

(iv)

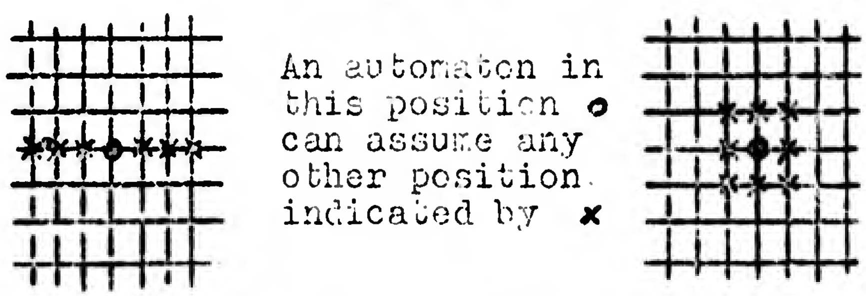

A o A o A A o B

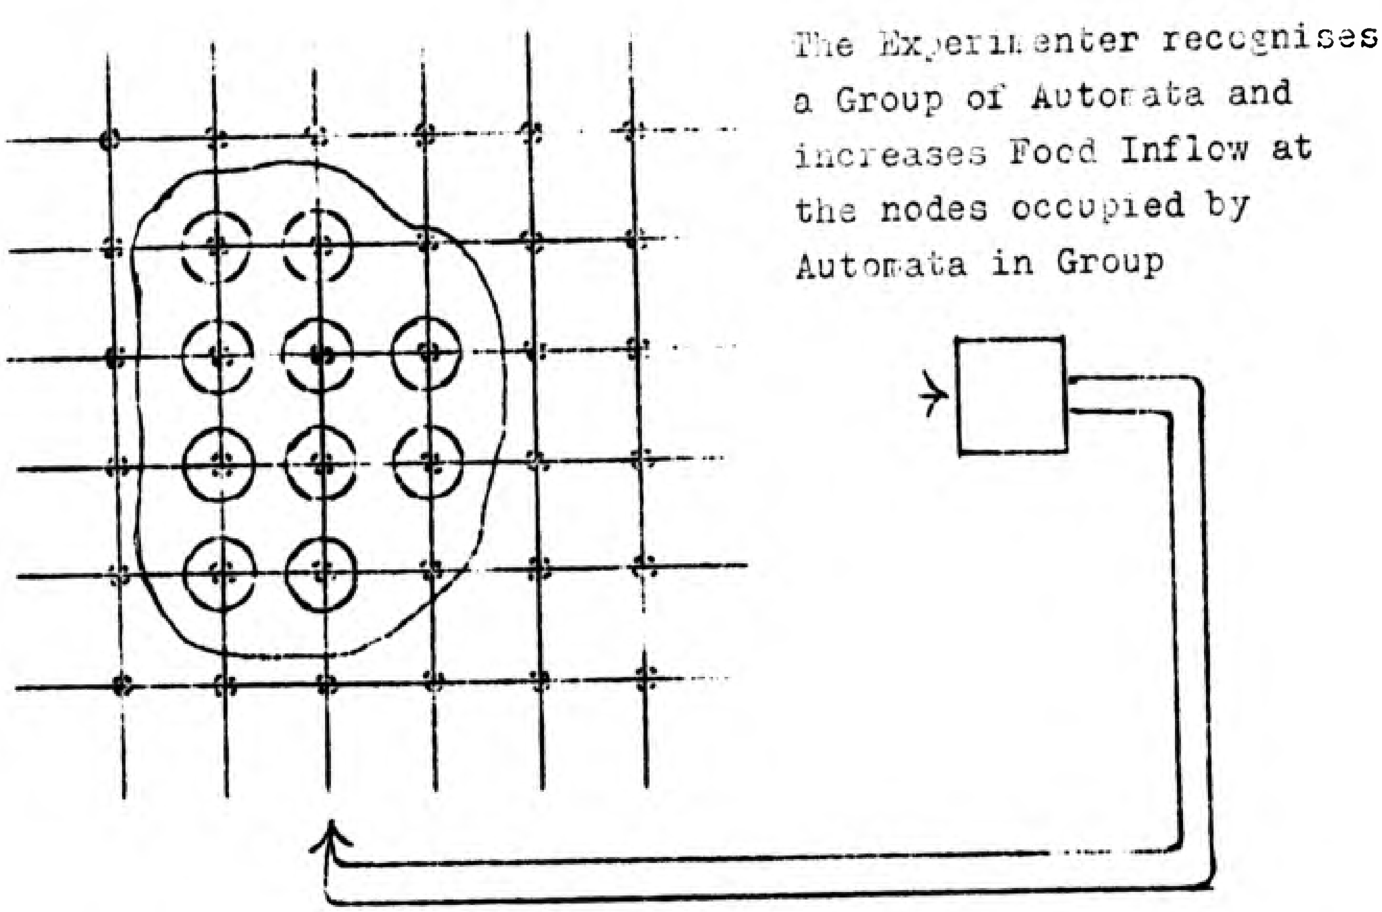

DIAGRAM 5.

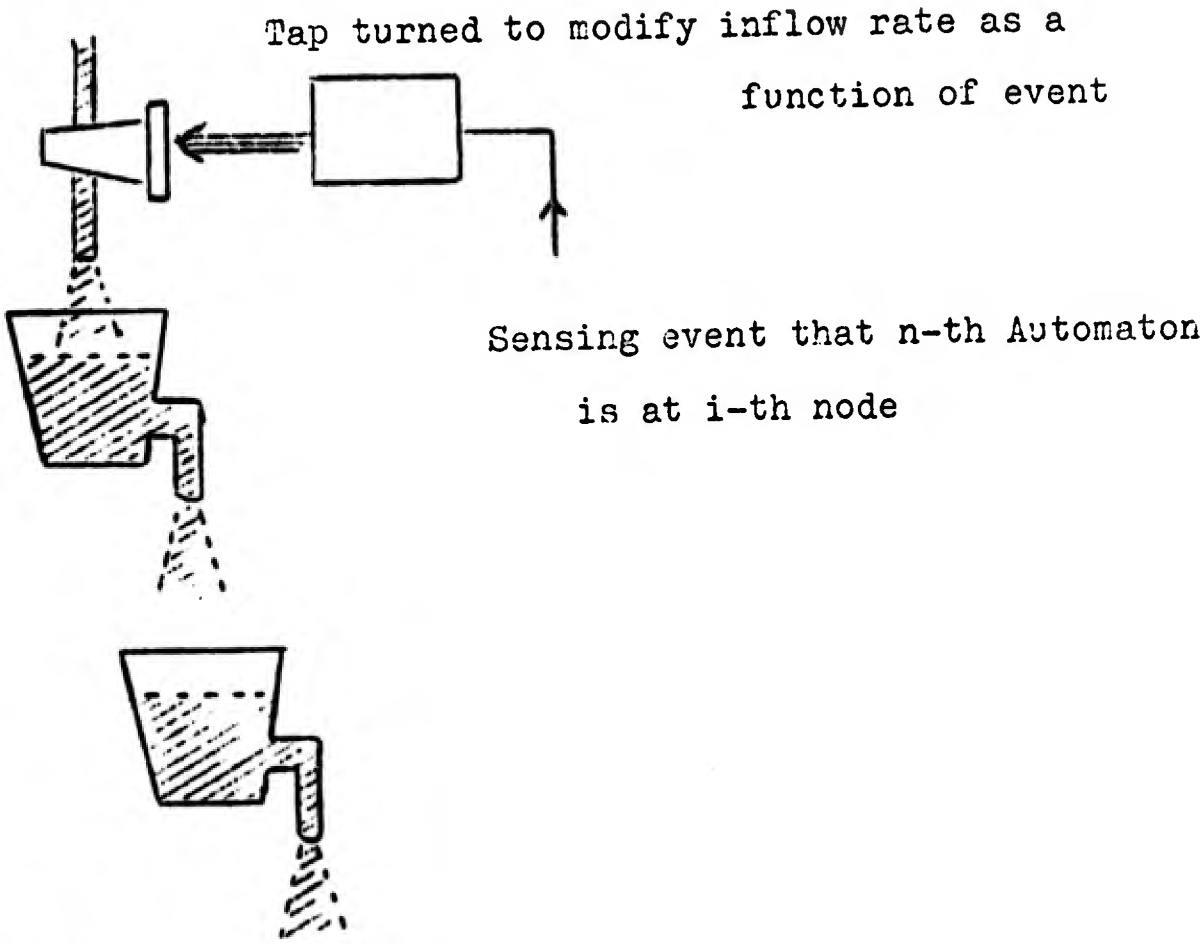

DIAGRAM 6.

DIAGRAM 7.

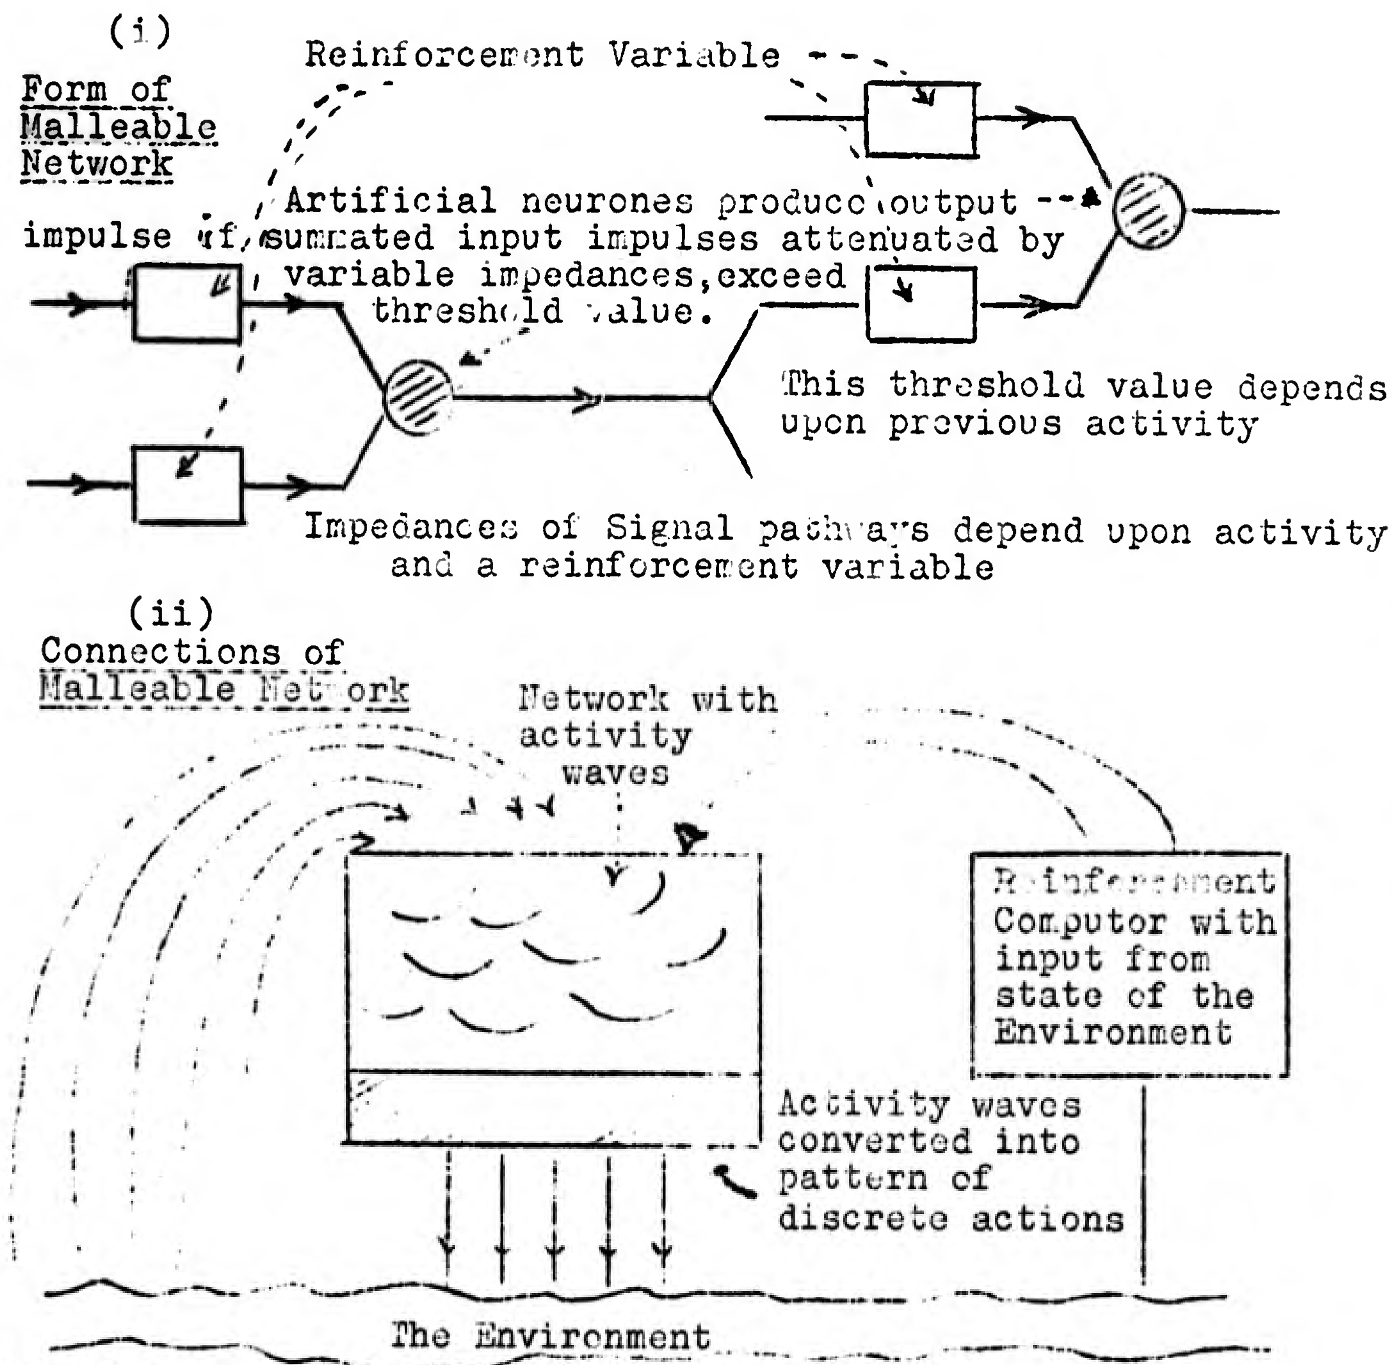

DIAGRAM 8.

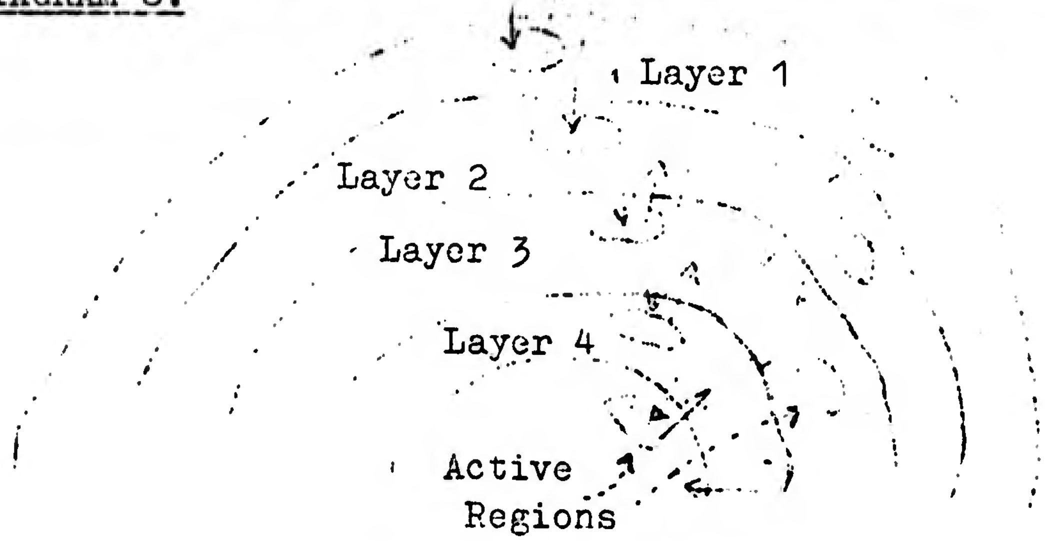

REFERENCES

1. W.R. Ashby. An Introduction to Cybernetics, Chapman and Hall, 1956.

2. C. Shannon, W. Weaver, Mathematical Theory of Communication, University of Illinois, 1948.

3. Stafford Beer, Towards a Cybernetic Factory, Symposium on Principles of Self Organisation, Urbana 1960. To be published. Ed: Heinz von Foerster and George Zoff.

4. Heinz von Foerster, The Self Organising System and its Environment, Chicago Symposium on Self Organising Systems, 1959. Pergamon Press 1960.

5. D.M. MacKay. Quantal Aspects of Scientific Information, Phil. Mag. 1950, 41, 289.

6. Gordon Pask, Physical Analogues for the Growth of a Concept. Proc. N.P.L. Symposium on Mechanisation of the Thinking Process. H.M.S.O. 1959.

7. Gordon Pask, Growth Process in a Cybernetic Machine, Proc. 2nd Congress. Int. Asscn. Cybernetics, 1958.

8. Heinz von Foerster, Gordon Pask, Predictive Model For a Self Organising System, Cybernetica Vol 1., 1961, 4, 1960.

9. Lettvin, Matturana, McCulloch, Pitts, What the Frog's Eye Tells the Frog's Brain, Proc. I.R.E. Nov. 1959.

10. J. Piaget, The Construction of Reality in the Child, Basic Books, 1954.

11. N. Rashevsky, Mathematical Biophysics, Dover 1960.

12. R. Feldman, A Homogeneous Network Approach to Self Structuring Systems, Microwave Research Institute, Brooklyn Polytechnic Inst., Brooklyn 1961.

13. W. Ross Ashby, Design for an Intelligence Amplifier, Automata Studies, Edited, Shannon & McCarthy, Princeton 1954.

14. Warren McCulloch, Agatha Tyche, Proc. N.P.L. Symposium on Mechanisation of Thought Processes, H.M.S.O., 1959.

15. Tyler Bonner, The Evolution of Development, Cambridge University, 1960.

16. Stafford Beer, The Irrelevance of Automata, Proc. 2nd Congress Int. Asscn. Cybernetics, 1958.

17. Georges Boulanger, Presidential Address, Congress. Int. Assocn. Cybernetics, 1956.

18. Contributions by Author to Automated Teaching and Programmed Learning, Edited by Lumsdaine and Glaser American Asscn. for Education, 1960.

19. W.S. McCulloch, W.R. Pitts, how we know Universals: the Perception of Auditory and Visual Forms, Bulletin Math. Biophysics. 9, 1947.

20. Heinz von Foerster, et al. Preorganisation for a Self Organising System, University of Illinois, Report on Contract Work, 1834, 21.

21. A.M. Uttley, Conditional Probability Computing in a Nervous System, Proc. N.P.L. Symposium, Mechanisation of the Thinking Process, H.M.S.O. 1959.

22. Frank Rosenblatt, Theorems on Statistical Separability in a Perceptron, Proc. N.P.L. Symposium on Mechanisation of the Thinking Process., H.M.S.O. 1959.

23. Alexi Ivanhenko, Royal Society Visit, Lectures and Discussion. 1960.

24. R.L. Beurle, Properties of a Block of Cells Capable of Regenerating Impulses, Proc. I.E.E., 240 B, 1956.

25. Farley and Clarke, Activity in Network of Neurone Like Elements, Proc. 4th London Symposium, Information Theory, To be Published. Edited: Colin Cherry.

26. J.W.S. Pringle, Parallel Between Learning and Evolution, General Systems, Vol 1. 1956.

27. D.M. MacKay, The Epistemological Problem for Automata, Automata Studies, Edited: Shannon & McCarthy, Princeton 1954.

28. J.D. Crane, Neurister Studies, Stanford Research Inst. 1960.

29. Marvin Minsky and Oliver Selfridge, Random Nets, Proc. 4th London Symposium, Information Theory edited: Colin Cherry, to be published.

30. Oliver Selfridge, Pandemonium - A Paradigm of Learning Proc. N.P.L. Symposium on Mechanisation of the Thinking Process H.M.S.O. 1959.

31. G.R. Bishop, Environmental Feedback in Brain Function Proc. Chicago Symposium on Self Organising Systems, edited: Yovits and Cameron, Pergamon Press 1960.

32. D.M. MacKay, Many papers, of which the most recent was presented at the 1961 Conference of the British Society for Philosophy of Science.

33. A.W. Burke Computation, Behaviour and Structure, Proc. Chicago Symposium on Self Organising Systems, Edited: Yovits and Cameron, Pergamon Press 1960.

34. S.T. Bok, A Vacuole Theory of Memory, Paper read at this Congress of Int. Asscn. Cybernetics, Namur, 1961.

Reference is also made to relevant points developed more fully in Gordon Pask, An Approach to Cybernetics, Hutchinson 1961, and the author's contribution to the Urbana Sympsium on the Principles of Self Organisation, where the evolutionary model is exhibited in greater detail.