

— AFOSR 66-0644 —

SYSTEM RESEARCH LTD

ANNUAL SUMMARY REPORT NO. 3.

Contract AF 61(052)-640. 1st Sept.62 - 31 Aug.65

Gordon Pask. M. Elstob. G. L. Mallen.

CLEARINGHOUSE
FOR FEDERAL SCIENTIFIC AND
TECHNICAL INFORMATION

Hardcopy	Microfiche	
\$4.00	\$0.75	104 pp 02

ARCHIVE COPY

Code 1

20 HILL RISE
RICHMOND
SURREY
ENGLAND
RIC 0801

AD631634

DOCUMENT CONTROL DATA - R&D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) System Research, Ltd. Richmond, Surrey, England		2A. REPORT SECURITY CLASSIFICATION <input checked="" type="checkbox"/> Unclassified Other - Specify	
3. REPORT TITLE Research on Cybernetic Investigation of Learning and Perception, Annual Summary Report No. 3.		2B. GROUP	
4. DESCRIPTIVE NOTES (Type of report and inclusive dates) <input checked="" type="checkbox"/> Scientific Report <input type="checkbox"/> Final Report <input type="checkbox"/> Journal Article <input type="checkbox"/> Proceedings <input type="checkbox"/> Book			
5. AUTHOR(S) (Last name, first name, initial) Pask, Gordon; Elstob, M.; and Mallen, George L.			
6. REPORT DATE AS PRINTED Feb. 66	7A. TOTAL NO. OF PAGES 70	7B. NO. OF REFS 7	
8A. CONTRACT OR GRANT NO. AF 61(052)-640	9A. ORIGINATOR'S REPORT NUMBER(S) (if given)		
B. PROJECT NO. 9769	9B. OTHER REPORT NO.(S) (Any other numbers that may be assigned this report) AFOSR 66-0644 AD		
C. 61445014			
D.			
10. AVAILABILITY/LIMITATION NOTICES Distribution of this document is unlimited.		<input checked="" type="checkbox"/> Available from DDC <input checked="" type="checkbox"/> Available from CFSTI <input type="checkbox"/> Available from Source <input type="checkbox"/> Available Commercially	
11. SUPPLEMENTARY NOTES (Citation)		12. SPONSORING MILITARY ACTIVITY AF Office of Scientific Research (SRI) Office of Aerospace Research Washington, D. C. 20333	
13 ABSTRACT Research is summarized on models that describe the learning of a structured skill and on simulations of populations of automata that become more complex as they develop. Applicability and limitations on a simple learning model based on terms of continuous, information-like measures are discussed. The model considers the contribution from learning of the i-th skill to learning of the j-th. Limitations arise for the description of learning of higher-order concepts. The relevance of statistical and homeostatic approaches to the description of learning and adaptation is considered; each is viewed as contributing to the characterization of a real-life population of organisms. The simulation model shows that individual automata do not learn on their own but in cooperating groups. The elaborate population that is postulated shows stability over a larger range of cost parameter values in an unconstrained environment than in a constrained environment. A gregarious automaton is described that has a sensory system (sensitivity to density of population) and a memory system; significance is associated with properties that remain invariant or exhibit regular and correlated transformation. Two appendixes are included that consist of preliminary drafts of two chapters of a manuscript on "The Cybernetics of Living Systems" and are entitled "Cybernetic Models and Control Systems" and "Fabric and Organization of Memory."			

Contract No:
AF 61(052)-640

15th February, 1966.

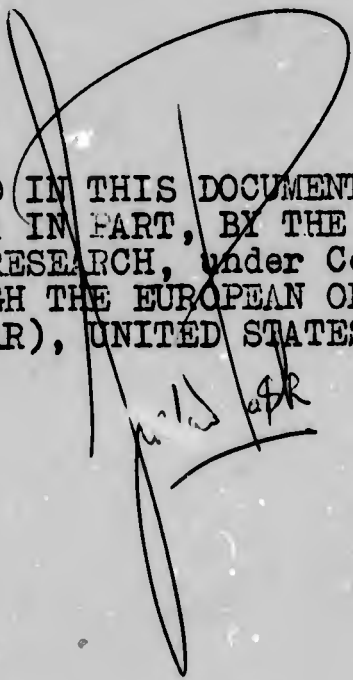
ANNUAL SUMMARY REPORT NO.3.

Research on Cybernetic
Investigation of Learning and Perception.

Gordon Pask
M. Elstob
G.L.Mallen

System Research Ltd.
20, Hill Rise,
Richmond, Surrey,
England.

"THE RESEARCH REPORTED IN THIS DOCUMENT HAS BEEN SPONSORED IN WHOLE, OR IN PART, BY THE AIR FORCE OFFICE OF SCIENTIFIC RESEARCH, under Contract AF 61(052)-640, THROUGH THE EUROPEAN OFFICE OF AEROSPACE RESEARCH (OAR), UNITED STATES AIR FORCE".



MONITORING AGENCY DOCUMENT NO:

ASTIA DOCUMENT NO:

Contract No. AF 61(052)-640

ANNUAL SUMMARY REPORT NO.3.

Research on Cybernetic
Investigation of Learning and Perception

Contractor: System Research Ltd.,
20, Hill Rise,
Richmond, Surrey,
England.

Project Staff:

Principal Investigator Gordon Pask

Research Associates B.N. Lewis, G.L. Mallen.

Research Assistants C.H. Moore, D. Watts, A. Shackleton.

Administration E. Pask.

Secretarial C. Merlin, V. Johnson.

15th February, 1966.

LIST OF CONTENTS

	<u>Page</u>
1. Introduction	1
2. <u>Part I.</u>	2
<u>Analysis of Cybernetic Learning Model.</u>	
Mathematical Model for the Learning of a Structured Skill	3
Subject Assumptions	4
The Model	5
Extension of the Model	7
Positive Transfer of Training	8
Stability of the Coupled System	9
3. <u>Part II</u>	12
<u>Automaton Models and Simulated Populations</u>	
Programme Development	12
Implicit Mechanisms	14
Explicit Mechanisms	17
Social Norms	21
Main Results	22
References	26
4. <u>Part III</u>	27
<u>Discussion of Criteria for Designing Additional Programmes</u>	
5. <u>Part IV</u>	30
<u>Further Developments of the Cybernetic Learning Model</u>	
6. <u>Appendix I</u>	
<u>Chapter II of "The Cybernetics of Living Things",</u> Cybernetic Models and Control Systems	
7. <u>Appendix II</u>	
<u>Chapter IV of "The Cybernetics of Living Things",</u> Fabric and Organisation of Memory	

Introduction.

The work carried out during the year can be conveniently classified under four headings of which a pair have already featured in the Administrative Reports No.8 and No.9. These headings are:

- (1) Approximate mathematical models for adaptively stabilised man machine interaction (and analysis of data obtained from experimental systems).
- (2) Final experiments with the existing form of automaton population computer programme.
- (3) Discussion and formulation of a novel class of programmes (to comprehend salient features of the automaton population and the algorithmic programmes).
- (4) Development of the basic Cybernetic Learning Model to include processes such as Imprinting and Maturation or their mechanical analogues.

The work under (1) has been carried out by G.L.Mallen and I and, so far, has led to a single publication describing a linear approximation model which is shown to be Lyapunov stable and in accord with the experimental data. So far as (2) is concerned, T.R.McKinnon Wood and I are chiefly responsible. For the discussions involved in (3) we had the benefit of Dr. B. van Hoijer's participation. Dr. Hoijer visited the laboratory for a couple of months in the summer; Mallen, McKinnon Wood and I also participated. Item (4) is my own effort; but currently the imprinting model is being simulated and some preliminary data should be available fairly soon.

Part 1. Analysis of the Cybernetic Learning Model.

Section (3) of Annual Summary Report No. 2. Contract AF.61(052)640 describes an investigation of some macroscopic properties of the basic learning model. Administrative Report No. 7. reports an experiment concerned with the adaptive instruction of a perceptual discrimination skill with eight subskills. Briefly, this experiment involves an 8 x 8 display of squares, each of which can be back illuminated. The track of an object is displayed by illuminating, in succession, a sequence of squares whilst groups of squares are cyclically illuminated to provide a "random" background. The skill consists in identifying and intercepting the displayed trajectories as soon as possible after their commencement. The difficulty of discriminating a trajectory is varied by changing the intensity of the random light pattern. An Adaptive Control Mechanism (A.C.M.); selects trajectories from a class of $N(n) = 1, 2, \dots, 8$ members at the n th. trial. The (A.C.M.), adjusts the difficulty levels, of the trajectories, so that the subject's proficiency ρ remains constant at a value ρ_0 predetermined by the experimenter. Thus, if ρ_i is the subject's proficiency with the i th. trajectory, at some stage in the instruction, then the (A.C.M.) attempts to make $\rho_i - \rho_0 = 0$. The (A.C.M.) also changes the membership of $N(n)$ - from which it selects trajectories to present to the subject - according to the following rule:-

$$\begin{aligned} \Delta N(n) &= +1 & \text{if } \delta\rho^+ > \delta\rho^- & \text{ unless } N(n)=8 \text{ when } \Delta N(n)=0 \\ \Delta N(n) &= 0 & \text{if } \delta\rho^+ = \delta\rho^- & \\ \Delta N(n) &= -1 & \text{if } \delta\rho^+ < \delta\rho^- & \text{ unless } N(n)=1 \text{ when } \Delta N(n)=0 \end{aligned}$$

Where $\delta\rho^+$ and $\delta\rho^-$ are the estimated values of $\lambda\rho/\lambda t$ (i.e. rate of learning) for $N(n)$ and $N(n)-1$ respectively.

The results of the experiments with this arrangement suggest certain conservation principles which were discussed in Administrative Report No. 7. These principles lead to a couple of analytic procedures, namely:-

- (1) We may regard the coupled man/machine system as Liapunov stable.
- (2) We may regard the subsystems concerned with the η_1 adaptive stabilisation of the subskills as quasi-independent components in the sense that compared to the relaxation time of the 'N' adjustment system, any one of these subsystems has a short relaxation time. Hence the entire system can be partitioned.

A Mathematical Model for the Learning of a Structured Skill.

This work is fully reported in a paper by Pask and Mallen(1). It is a mathematical description of the basic system behaviour when a subject learns a structured skill under adaptively controlled conditions. It may also be argued that the model exhibits features which are common to more general cases of learning. According to the familiar cybernetic learning model, as expounded in Annual Summary Report No.2. and elsewhere(2) the acquisition of a subskill is correlated with the development (in the subject) of an organisation that realises and embodies the subject's codes for this subskill. In the mathematical model presented here, the changes that occur when the subject learns under adaptively controlled conditions, are represented by the changes in a pair of information measures, namely H and I. Where H is the subject's uncertainty and I is the subject's maximum possible uncertainty and is determined by the experimental environment.

We envisage, in the first instance, that our model describes the acquisition of a single subskill and that the experimental environment has the following general features.

A subject is presented with a sequence of problem stimuli at a steady rate through-out a sampling interval Δt . The sequence has an intrinsic uncertainty I associated with it. At the end of the sampling interval an adaptive control mechanism evaluates the subject's responses, and uses this information, according to some control rule, to adjust the content of the next sequence of problem

stimuli. If the problem stimuli represent components of a structured skill then the ACM can control and even optimise subject learning. Experiments conducted along these lines were reported in Annual Summary Report No.2.

Subject Assumptions.

Certain assumptions are made about the conduct of the subject in the experiment.

- (1) We assume that the subject attends to the field of attention determined by the experimenter and interacts within the experimental environment "according to the rules of the game".
- (2) We assume that the subject prefers to achieve goal satisfying responses rather than responses that do not satisfy the accepted goals.

Accepting these axioms, the following statements can be made regarding the abilities and behaviour of the subject.

- (1) The subject has an upper and lower limit to his rate of mentation; if he is presented with material outside these limits he is unable to cope with it. We can express this condition in our descriptive model by stating that for learning to occur, the subject's uncertainty H , must at all times be such that

$$H_{\max} > H > H_{\min}$$

Where H_{\max} and H_{\min} are the maximum and minimum uncertainties that the subject is able to cope with.

- (2) Learning will occur (i.e. $\frac{dH}{dt} < 0$) if

$$H_{\max} > H > H_{\min}$$

- (3) We assume that for all trials prior to reaching criterion performance (in a subskill)

$$\frac{dH}{dt} > 0 \text{ if } H_{\max} > H > H_{\min}$$

- (4) We assume that the subject acts as a self organising system in the sense of von Foerster, i.e. the rate of change of behavioural redundancy is positive.

The model that we now present for a single subskill is the simplest that accounts for these conditions.

The Model.

For learning to occur, we know that, $H_{\max} > H > H_{\min}$.
Let H_0 represent the subject's uncertainty at the

beginning of a trial. Then for learning to occur in that trial,

$$H_{\max} > H_0 > H_{\min}.$$

We assume that the subject's learning rate ($-\frac{dH}{dt}$),

is proportional to :-

- (1) A subject's characteristic parameter, k .
- (2) The amount by which his initial uncertainty H_0 exceeds H_{\min} , i.e. ($H_0 - H_{\min}$).
- (3) His remaining uncertainty H . Combining these, we may write:-

$$\frac{dH}{dt} = -k (H_0 - H_{\min}) H \quad \text{Hence}$$

$$H = H_0 \exp(-k(H_0 - H_{\min})t) \dots(1)$$

It is reasonable to assume that $\frac{\Delta H}{\Delta I} = \text{constant}$.

This simply implies that the subject's uncertainty is related, in a linear fashion, to the intrinsic uncertainty of the experimental environment. For simplicity of description we assume:-

$$\Delta H = \Delta I \dots(2)$$

Next, we express the general form of the control action taken by the adaptive control mechanism. The general rule for a single subskill can be written as follows:-

$$I_n \Delta t = I_{(n-1)} \Delta t + \alpha \Delta I \quad \text{if } H_{n \Delta t} > H_{(n-1) \Delta t} \dots(3a)$$

$$I_n \Delta t = I_{(n-1)} \Delta t - \beta \Delta I \quad \text{if } H_{n \Delta t} < H_{(n-1) \Delta t} \dots(3b)$$

$I_{n \Delta t}$ is the value of I at the n -th trial, (trial length is Δt).

$H_{n \Delta t}$ is the value of H at the n -th trial, before control action is taken.

α and β are control parameters.

Considering now the behaviour of this basic model.

$$\text{At } t = 0 \text{ let } I = H_{t_0} = H_{\min}$$

then one trial later,

$$H_{t_0 + \Delta t} = H_{t_0}, \text{ from (1)}$$

And the control action to be taken is determined by (3a); that is,

$$I_{t_0 + \Delta t} = I_{t_0} + \alpha \Delta I.$$

and, $H_{t_0 + \Delta t}$ is increased to $H_{t_0} + \alpha \Delta I$.

During the period Δt to $2\Delta t$ $H_o > H_{min}$ and learning can and does occur. The subject's uncertainty at the end of this period is

$$H_{t_0+2\Delta t} = H_o(t_0 + \Delta t) \exp(-k \alpha \Delta I \Delta t)$$

The decrease in H due to learning in that trial is $-\Delta H(t_0+\Delta t) = H_o(t_0 + \Delta t) (1 - \exp(-k H_o(t_0 + \Delta t) - H_{min}))$,

and in general the reduction in H in trial t is $-\Delta H(t) = H_o(t) (1 - \exp(-k(H_o(t) - H_{min}) \Delta t)) \dots (4)$

We can see that $\Delta H(t)$ increases as t increases. At a time t_1 , say, a point will be reached when

$\Delta H(t_1)$ is exactly equal to the increase in I ,

namely ΔI introduced by the A.C.M. For times after t_1 the initial uncertainty at the commencement of each trial (i.e. $H_o(t)$) will be $H_o(t_1)$. This

will continue to be the case until $I = I_{max}$ at time $t = t_2$.

If the terminal criterion for a subskill is that $I_{T \Delta t} = I_{max}$ (T = no. of trials to reach criterion performance) and $H_o(T \Delta t) = H_{min}$. Then the time taken to reach this terminal condition will be $T \Delta t = t_2 + t_d$.

Where t_d = time required for H to decay from $H_o(t_2)$ to H_{min} at constant $I = I_{max}$. From (1) this time

is $t_d = \frac{\log H_o(t_2) - \log H_{min}}{k(H_o(t_2) - H_{min})}$.

Thus, for a specific $H_o(t_2)$, t_d is dependent only on subject parameters and so is constant. Hence, $T \Delta t$ is minimised if $t_2 - t_1$ is minimised. $t_2 - t_1$ will be minimum when $H_o(t_1)$ is maximised and from

(4) we see that this occurs when $H_o(t_1) = H_{max}$.

Then $\Delta H_o(t_1)_{op} = \Delta I_{cp} = H_{max} [1 - \exp(-k(H_{max} - H_{min}) \Delta t)] \dots (5)$.

ΔI_{op} is the optimum value of the discrete "difficulty" increase to give maximum learning rate. At equilibrium (after t_1 , when $\frac{dH}{dt} = 0$).

$$\frac{dI}{dt} = \frac{\alpha \Delta I}{\Delta t}$$

Now, taking the linear approximation to equation(1) :-

$$H = -k(H_0 - H_{\min})t.$$

$$\frac{dH}{dt} = -k(H_0 - H_{\min}).$$

we note that $\frac{dI}{dt} = -\frac{dH}{dt}$ at equilibrium

and since $\frac{dH_0(t)}{dt} = 0$ we may put

$$\frac{dH_0(t)}{dt} = \frac{dI}{dt} + \frac{dH}{dt} \dots\dots\dots(6)$$

$$i. e. \quad \frac{dH_0(t)}{dt} = \frac{dI}{dt} - k(H_0 - H_{\min}) \dots\dots(7)$$

These equations describe the behaviour of the coupled subject/machine system. $\frac{dI}{dt}$ is determined by the

A.C.M control rule.

We can interpret equation (6) as follows:

If due to some external interference, the subject begins to learn more quickly at $H_0(t) = H_{\max}$ then $\frac{dH}{dt}$ goes more negative, and hence $\frac{dH_0}{dt}$ goes negative.

This is recognised by the A.C.M. which then increases $\frac{dI}{dt}$ to compensate. For a single subskill the A.C.M. achieves this by changing α , the control parameter,

Extension of the Model.

So far the model has been described in terms of the continuous, information-like measures I and H, These are not easily measured during the learning process, but it is argued that they can be replaced in the relevant equations, by the sampled state variables, difficulty η_i and proficiency ρ_i .

In particular we can substitute

$$\Delta \rho \text{ for } -\frac{dH}{dt} \text{ and } \Delta \eta \text{ for } \frac{dI}{dt}$$

The model describes the processes that take place in the dynamic man/machine system when a single subskill alone is being learnt. In order to

extend the description to include the case when several subskills are rehearsed jointly, we must consider the cases of positive transfer of training and conditional interference among the various subskills. The single subskill model is characterized by the time constant Δt of the trial period or sampling interval. The A.C.M. takes control action at the end of periods of this length. In the multiple subskill case, or the case of a structured skill, control action for a particular subskill may be much delayed, since the time between one rehearsal and the next, may be relatively long. Thus the control mechanism must have a much longer time constant. The subsystem parameters change very slowly compared to the subsystem time constants Δt and so, approximately, we may regard the subsystems as separable, and if the complete system is stable, in dynamic equilibrium.

Positive Transfer of Training.

Suppose there are 'm' subskills and that we can attach 'a priori', values a, b, \dots to the interactions among them. That is we are given an interaction matrix

$$Z = (Q_{ij}) \text{ where } Q_{ij} = 0 ; i = j \\ Q_{ij} > 0 ; i \neq j.$$

The elements of this matrix indicate the amount by which the achieved state I_i of the i -th subskill aids learning rate in the j -th subskill. Consider the case when $m = 2$.

$$Z = \begin{bmatrix} 0 & a \\ b & 0 \end{bmatrix} \quad a, b, > 0.$$

Let us assume the subject has been brought to equilibrium, by the A.C.M. on each of the subskills individually before joint rehearsal takes place. Let t_1 be the earliest time at which both subskills are at equilibrium with $H_0 = H_{\max}$. Let t_2 be the earliest time at which $I_i = (I_{\max})_i$ for $i = 1, 2$.

Then from equation (7) we get two coupled equations specifying the new dynamic conditions for the equilibrium of the two subsystems. These are

$$\frac{dI_1}{dt} = k_1((H_{\max})_1 - H_{\min}) + b I_2 \dots\dots\dots(8a).$$

$$\frac{dI_2}{dt} = k_2((H_{\max})_2 - H_{\min}) + a I_1 \dots\dots\dots(8b).$$

now t_2 is determined by the lowest average $\frac{dI_1}{dt}$

and it is evident under the joint rehearsal conditions that this is minimised if $\frac{dI_1}{dt} = \frac{dI_2}{dt}$;

assuming that $k_1 = k_2$ and $(H_{\max})_1 = (H_{\max})_2$.

If we wish to maximise the expected rate of learning we may use the control policy indicated by this result. The A.C.M. can achieve this control by altering the microscopic difficulty levels ($\propto \Delta I_1$), and ($\propto \Delta I_2$) of the two subskills or by varying the relative number of times each subskill is rehearsed.

The above analysis uses a linear approximation and it becomes inapplicable at the point, which is realised in real learning, at which the subject is using a higher order concept. When he has acquired this higher order concept the two subskills are treated as a whole and interaction between the two steps.

The case of conditional interference can be treated in a similar manner. The same rule optimising learning rate applies. Again, the linear approximation breaks down when the subject acquires a higher order concept.

Stability of the Coupled System.

Returning to the question of stability as exhibited by the coupled man/machine system. It is not altogether surprising that the coupled system shows features of stability since it is the purpose of an Adaptive Teaching Mechanism to make adjustments in the teaching environment to compensate for the changes that occur as a consequence of the subject's learning.

We will show below that it is possible to apply Lyapunov's criteria for dynamic stability to the coupled system. We can obtain a Lyapunov function for the system in terms of the behavioural redundancy R , of the subject. This will enable us to assert that a condition for stability is that the subject be able to act as a self-organizing system.

According to Lyapunov's theorem, any dynamic system with state variables z_i is stable if V is a function of these variables such that:-

$$(1) \quad 0 > dV/dt. = \sum_i dz_i/dt \cdot \partial V/\partial z_i$$

$$(2) \quad V > 0.$$

The state variables of the coupled man/machine system are the $2n$ values of ρ_i and η_i . Since, on the average, $\rho_i = \rho_0$ for all "i" (due to A.C.M. activity), we can state that $d\rho_i/dt = 0$. Then, the first Lyapunov stability condition (for the coupled system) becomes:-

$$0 > dV/dt. = \sum_i d\eta_i/dt \cdot \partial V/\partial \eta_i.$$

We argue that there is a function of η_i , and hence of I , which has the properties of a Lyapunov function for the system. Let this function be $f(I)$, then we require that:-

$$(1) \quad 0 > df(I)/dt.$$

$$(2) \quad f(I) > 0.$$

Choose $f(I) = H/I = 1-R$ where R is the behavioural redundancy. Then:-

$$df(I)/dt. = 1/I \cdot dH/dt. - H/I^2 \cdot dI/dt.$$

Now, when the coupled system is in equilibrium we have shown that $dH/dt = 0$. Thus:-

$$df(I)/dt. = - H/I^2 \cdot dI/dt.$$

Hence

$$0 > df(I)/dt. \quad \text{iff} \quad dI/dt > 0.$$

Thus the system is stable if, and only if, $dI/dt > 0$. In subject assumption 4, we assumed that, for learning to occur, the subject acts as a self-organizing system i.e. that $dR/dt > 0$ where $R = 1-H/I$. This condition implies that $dI/dt > 0$. Hence, the condition for $f(I) = 1-R$ to be a Lyapunov function is the condition that the subject be a self-organizing system.

References:

- (1) Pask G. and Mallen G.L. - The Method of Adaptively Controlled Psychological Experiments. Proc. of I.F.A.C. Symposium, Teddington 1965. Instrument Society of America, 1966.
- (2) Pask G., - Man as a system that needs to learn. Advances in Cybernetics. S. Beer, F.H. George and D. Stewart (eds.). New York Academic Press. In press.

BLANK PAGE

Part 2. Automaton Models and Simulated Populations.

(1). Programme Development.

In Administrative Report Number 9, we presented fairly detailed records (automaton distributions at moves 5,10,25,55, and 100), for a couple of simulations that differed only insofar as in (I) there were environment constraints intended to favour the complex types "x" and "+" whilst in (II) these constraints were deleted. The constraints prohibited simple automata (those of types "-" and "1") access to certain nodes which were thus exclusively the feeding ground for complex automata.

It turned out that the complex automata were dramatically dominant in (I) and in (II). In each case the simple automata became extinct as might be expected. The rate of extinction was rather greater in (I) than in (II) and the steady state population number of complex automata is greater in (I) than in (II).

This sort of dominance does, of course, depend upon an initially diffuse distribution of automata (the main experimental condition for Administrative Report Number 9). Entirely different patterns (in which the simple types of automata are relatively more numerous) can develop from compactly structured initial distributions, for example, the populations in Graph 1.I and in Graph 1.II, of the induction experiments described in Annual Summary Report Number 2. However, there is a sense in which a "simple dominance" is metastable; transitions "1 or - + or X" are permitted but the converse transitions "t or x and 1 or -" are not. The induction effects produced in the experiments of Annual Summary Report Number 2, all depend, in one way or another upon this effect; similarly, of course, biological induction depends upon the metastability of equipotential cells which, in this case, is due to the probability of differentiation and the improbability and usually pathological occurrence of dedifferentiation.

To ensure that the complex dominance effect in (I) and in (II) is not an accident, these simulations have been continued for a further sequence of 100 moves (making 200 moves in all of which the first 100 have already been reported). The results appear in Graph 1. and in Graph 2 of the present Report.

We now consider reducing the dominance of these complex types of automata by selectively increasing the cost of structural maintenance for all of "x" or "+", and decreasing it for "-" or "1".

It is fairly easy to find a value of the structural maintenance cost parameter (the amount of food in the stomach of an automaton that must be expended, each move, to maintain the automaton body at a given age level) which renders the steady state population instable and which will render it extinct because automata that are able to reproduce fail to make the encounters that are a prerequisite for reproduction before they decay. The determination of this cost value is shown in Graph 3, for the environment with constraints and in Graph 4, for the environment without constraints. The fact that the critical value is higher in Graph 3 than in Graph 4 is due to the "filtering" effect of the environment constraints which tends to separate "x" from "+" into different habitats as well as favouring "x" and "+" against the "-" and "1" (the separation of habitats will, on average, provide the automata with more food).

Now we insert a more than normal but slightly less than critical value of the cost value into the initial conditions of a simulation (so that it acts upon the "x" and "+" automata before the "-" and "1" automata have become extinct). In this condition the maintenance of the "x" and "+" population numbers depends jointly upon the encounter of viable "x" or "+" and the encounter of combinations of viable simple automata, namely generative encounters of the type "-" with "1" into "x" or "+". Insofar as the complex population is marginally stable on its own it depends upon generative encounters between simple automata. Insofar as the larger than normal maintenance cost value reduces the number of complex automata at a given instant, the simple automata are able to survive. The dependency, cited above, renders the population of simple automata somewhat analagous to a food supply to the complex automata and so the complex automaton population can be stable if and only if it allows its food supply to survive by limiting its ^{own} rate of reproduction.

In view of these comments the data in Graph 5, for the constrained environment and in Graph 6, for

The unconstrained environment is not altogether surprising. As the value of the cost of maintenance parameter is increased from Graph 5 to Graph 6 for (I) and from Graph 7 to Graph 8 for (II) an oscillatory interaction develops between the simple automata and the complex automata. The form of oscillation is more evident in Graph 9 where the "x" and "+" automata are lumped together as complex types and the "-" and "1" are lumped together as simple types (for the constrained environment) and in Graph 10 (where the same procedure is adopted for the unconstrained environment). It resembles the Volterra like predator prey interactions that have been modelled in a statistical fashion by Kerner (1), and there is every reason to believe that the underlying process is some primitive form of the predator prey interactions he and others have considered in real life populations. As the cost of maintenance value is further increased the oscillation frequency increases (and predictably, since it is a non linear system) their amplitude also increases. Rather suddenly the amplitude of the oscillation becomes so great that one of the negative excursions reaches 0 for the simple automata. Since (by definition) the complex automaton population is instable alone at this value of the cost parameter and since there are no simple automata left, the entire system decays.

There is no explicit rule embodied in the programme that allows the automaton population to restrict its rate of reproduction. Yet it is evident that a restriction is applied, for otherwise the oscillatory behaviour would be, at the most, of ephemeral occurrence. In fact, this oscillatory behaviour persists over a considerable range of values of the cost parameter.

(2) Implicit Mechanisms

The mechanism involved is of some interest because the commentators on Kerner's work (1), often appear to regard his statistical conclusions as contrary to the point of view adopted by Wynne Edwards (2) in his discussion of homeostatic mechanisms for the control of population numbers. In fact there is no necessary disparity between the statistical and the homeostatic point of view in this matter, and the present simulation, primitive though it is when compared with a real life population, serves as a vehicle for the unifying argument that will be briefly presented.

According to Wynne-Edwards (2) thesis (which was mentioned in Annual Summary Report 2) the local population in a given habitat either inherits or develops conventional constraints which act as a homeostatic mechanism for maintaining the number of members of the population near to an optimum value determined by the food available from this habitat. The sensory mechanism is usually a ritualised Epidiestic display, for example, the ritual singing of male territorial birds and their habit of flying around their territory. Such a display allows any individual to glean and to compare information regarding the fruitfulness of the habitat and its number of occupants. The operation that is performed by the homeostatic control system is a limitation of the rate of reproduction or alternatively an operation that leads some members of the population to emigrate into another habitat. (Several, rather varied, mechanisms have been cited as responsible). The decision regarding which organisms shall be prevented from reproducing or induced to emigrate is made by an hierarchical ordering of the members of the population which is analagous to, and for some species may be identical with, the familiar peck ordering of fowls. Something akin to an hierarchical ordering is built into the present simulation due to the manner in which automata with a high value of stomach content are processed before those with a lower stomach content. However, a similar ordering is built into any population wherein the members differ in vitality or speed of response (which is true of any realistic population) and where vitality or speed of response confers some advantage upon the organism concerned (this advantage may, of course, be overlaid by a conventional ordering). But the simulation contains no conventions that are analagous to the Epidiestic display or mechanisms analagous to a variation in reproduction rate. Wynne Edwards refers to the gamut of sensory activities involved in the Epidiestic display and the subsequent decision regarding which organisms shall be inhibited as "symbolic competition" (which is really a cooperative process) and the crucial issue is emphasised by the comment that our simulation programme makes no explicit provision for the cooperative process of "symbolic competition" (apart from the food procedence mentioned a moment ago).

In contrast, the idea of statistical interaction lacks any cooperative feature. The predator prey interaction is wholly competitive in character! As the predators eat up the prey, fewer predators are able to live in the habitat, hence the number of prey can increase and consequently more predators may, after some delay, be accommodated. If what amount to "weak interactions" between predators are subsumed by a "statistical mechanics" it is shown that the situation can lead to an oscillatory dynamic equilibrium; rather than the crass instability predictable on the basis of a deterministic approach.

Mathematically, this system is impeccable. From a practical point of view, however, we may profitably make a closer examination of the population of predators, and also of the weak interactions. For, on the one hand, it is evident that a homeostatic system would be an advantage if it existed. On the other hand, it is evident that almost any conceivable realised population will embody constraints and inertial terms that could provide an implicit homeostatic mechanism even in the complete absence of any manifest or explicit or predetermined arrangement. Such an implicit mechanism is readily adumbrated by weak interaction and in some cases exerts the required smoothing effect.

Our simulation, for example, provides an implicit homeostatic mechanism. The automata cannot learn on their own but cooperating groups of automata most certainly can learn. There is plenty of evidence that organised collections of automata act as sensory devices that detect the food concentration and possibly the automaton density in the neighbourhood of a cooperative group of automata. If, for example, the motile "x" automata are deleted from the replicating cooperative group of automata shown in DIAGRAM 8 of Annual Summary Report 2, the entire structure decays. It is likely that these automata have a sensory function with reference to the group and in less elaborate cases than DIAGRAM 8. such a function is perfectly evident. Similarly, even though there is no explicit mechanism for adjusting the rate of reproduction, there are many cooperative interactions that allow an aggregate of individuals to prevent the reproduction of an individual member of this aggregate or a particular type of member. These microscopic modes of interaction have been adequately reviewed in previous publications.(3)(4).

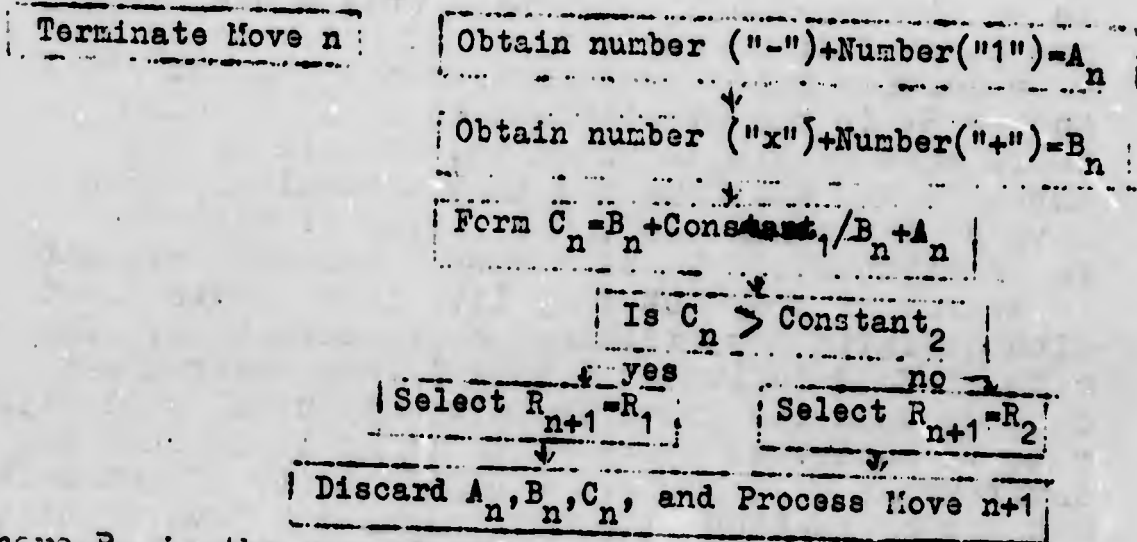
It is difficult, and it may be impossible, to hold a particular implicit mechanism as responsible for the manifest control of predator numbers (the homeostatic system is almost certainly a distributed control system). However, it can be argued that a greater number of microscopic features are available in some conditions than in others, for example, that the unconstrained environment provides a greater diversity of potential implicit mechanisms (whether or not these are in fact used) than the constrained environment. In this connection, it is worth commenting that the elaborate population is stable over a larger range of cost parameter values in the unconstrained environment than it is in the constrained environment.

The existence of implicit mechanisms able to form an implicit homeostatic system is a mechanical consequence of modelling, in a detailed fashion, the sort of organisation that is imaged at a macroscopic level by the competitive mathematical structure. These implicit mechanisms entail microscopic cooperation, often perhaps the mode of cooperation called symbolic competition. Hence, there is no hard and fast demarcation between a population that does and does not satisfy the paradigm of social homeostasis, although a ritualistic or conventional framework can be detected when this homeostasis is explicitly manifest. Since these comments are true for a rather crude model like the present simulation, they are presumably applicable in the case of any physical realisation, such as a real life population of organisms, where the constraints imposed by the physical fabric are more stringent and the inertial properties of this fabric are more elaborate. Thus, far from arguing that either thesis of mutualism is mistaken, I maintain that there is a continuum of which each thesis is one extreme case. In reality, the dynamics of mutualism are to be described by some point on the continuum. The only criticism is operational. Whilst the homeostatic mechanism remains implicit it may often be impossible to distinguish, according to operational criteria, whether or not it exists.

(3). Explicit Mechanisms.

It is possible to introduce explicit homeostatic mechanisms for the control of the number of automata

in the population at a macroscopic or at a microscopic level. The macroscopic control is mediated by the experimenter, but it could be embodied in an additional subroutine for the programme. The particular mechanism that has been adopted is not ideal and its form is chiefly determined by what is and is not an accessible part of the programme. It is easy, in practice, to obtain a print out of the number of simple automata "-" and "1", and the number of elaborate types "x" and "+". Further it is easy to modify the reproduction rules from move to move. Hence, the explicit homeostatic mechanism is merely:



Where R_n is the reproduction rule to be used in processing the automata at the n-th move and where R_1 is the reproduction rule used in the present simulation, namely;

R_1 = -o- = 75% -, 25% 1; 1o1 = 75% 1, 25% -;
 -o1 = 1o- = 50% "x", 50% "+" Xox = 100% X; + o += 100% +;
 and X o + = X . o + = 50% x, 50% +;

and R_2 is the reproduction rule

R_2 = -o- = 75% -, 25% 1; 1 o 1 = 75% 1, 25%; - o 1 = 1 o -
 = 50% x, 50% +; x o x = Nothing.

The effect of employing this macroscopic control system is exhibited in Diagram 11 and Diagram 12.

Explicit social homeostasis can also be introduced at a microscopic level to yield a distributed control system. The microscopic mechanism is a rule that is built into all or some of the automata and which, if suitably chosen, acts as a form of social constraint (the automata are designed to accept a limitation in the interest of the group even though it may be contrary to their own immediate welfare). Experiments with microscopic mechanisms have been performed with the earlier simulations, briefly mentioned in Administrative Report Number 5 that contained automata having 8 possible motions and imaged as $*$ in the usual convention (the $*$ automata are reached from simple types of automata by the reproduction rules in Administrative Report Number 5; in the simulations we shall discuss there were no simple automata, and all of the automata had the form $*\ast$; in the simulations we have so far considered, the $*$ type of automaton is prohibited). These 8 motion automata are not in search of prey but they live in an environment with definite "overfishing" characteristics. The rate of food influx at a node in the environment depends upon its previous occupancy until a limiting rate is achieved. Until this limit is reached the environment will support a high density of automata. But any collection of nodes is readily "overpopulated" for beyond this limit such a high density of automata cannot be supported. The population has "overfished" the environment and it cannot survive unless it is able to detect an approach to this critical value before the value is reached and to adjust its density. The fact that populations such as those imaged in Graph 13 and Graph 14 (for population number) and Graph 15 and Graph 16 (for the corresponding food average) were, in fact, stable is further evidence in favour of implicit mechanisms of social homeostasis, for without it overpopulation could hardly have been avoided in at least some regions (of course, overpopulation did occur, in many simulations).

Amongst several types of $*$ automata used in these experiments one was "gregarious". The i -th "gregarious" automaton contained an adaptive parameter assuming the value of P in at the n -th move.

It was equipped with a sensory arrangement for determining the density of automata (rather than the level of food) at the nodes to which it could move and at the n -th trial, moved to whichever node had a density of automata nearest to P_{in} . This sort of automaton also had a "memory" system wherein it could record the average increment of stomach food whilst P_i was being increased by +1 for each move, say Z_{in}^+ and the average value of this increment whilst P_i was being decreased by -1 for each move, say Z_{in}^- . It adjusted the value of P_i by making

$$\begin{aligned} \Delta P_{in} &= +1 \text{ if } Z_{in}^+ > Z_{in}^- \\ &= -1 \text{ if } Z_{in}^- > Z_{in}^+ \end{aligned}$$

this "adaptation" smoothed out the density distribution and stabilised some types of population, but it was liable, as we have already commented to an inherent instability that stems from the fact that aggregations of automata with high mean P_i values reproduce rapidly when they come across a favourable part of the environment but are unable to disperse when the food is depleted (they need a mechanism like the phase mechanism for locust emigration, but the programme does not provide it).

The gregarious automata were next developed into a form that used the automaton density data and the adaptively controlled value of P_{in} to adjust the critical stomach food level at which reproduction occurs, if there is a suitable encounter with another automaton (the motions of these automata were determined, as in the case of other non-gregarious automata, by the food levels at all accessible nodes. They move to the maximum food level available).

This arrangement is certainly an explicit homeostatic mechanism and its efficacy at the level of the population, rather than the individual automata, is indicated by comparison of the "normal" and the "homeostatic" curves in Graph 17 and in Graph 18. The perturbation applied to induce instability is a sudden, experimenter controlled decrease in the food level at all of the nodes in the environment. It is evident that the "homeostatic"

population is far more tolerant of this perturbation than the "normal" population from which it differs by the adjunction of the explicit control over rate of individual reproduction. This comparison is not exact because rather a long interval is needed (about 25 moves) for the P_{in} to assume representative values. However, the effect exhibited is dramatic and similar effects can be produced repeatably in different populations (there is a mass of partly analysed print out that supports this point on inspection).

(4) Social Norms.

There are many ways in which features of a developing population can be usefully described. In Annual Summary Report 1, for example, we presented detailed data about the motions of individual automata. It would have been possible to compute the correlations of motions and it might, in some conditions, be worthwhile. Less detailed information is provided by density distributions and move by move plots of population number or plots of the average numbers of different types of automata.

Amongst the more interesting sorts of data, is a curve representing the numbers of individual automata that possess a particular property. A couple of these curves are shown in Graph 19 and Graph 20. If these distributions are invariant over a reasonable number of moves or if they show regular and correlated transformations they are often more revealing than an account of the motions executed by the automata.

In sociological studies, for example, stable distributions of this sort (made with reference to the number of individual organisms having a sociologically cogent attribute value) are called social norms. Arguing conversely (and no longer restricting our attention to social attributes) the motion characteristics of automata in the environment constitute a very specialised and not necessarily very interesting set of norms. Indeed, suitable automata need not move in any such environment as a plane of nodes. They might equally "move" in a space of properties and the character of their motion might equally well be described by norms

or number distributions made with reference to any or all of these properties. This is the type of description and the type of automaton which is considered in Part 3 of the present report. The automata have an "abstract environment" which is, however, more readily identified with interesting facets of reality than a special plane of nodes.

(5). Main Results.

The present type of simulation has served a number of purposes and has been used to test several hypothesis. In some respects the picture it provides, though readily conceived at a superficial level, is cumbersome and of limited applicability. Hence, although the programmes will be preserved and may occasionally be used, it is not expedient to carry out further experimental studies in this framework. The picture remains, however, an admirable vehicle for demonstration.

In fact, the chief merit of such a simulation lies in the opportunity it affords for investigating mechanisms. The macroscopic characteristics of populations are often mathematically tractable and the mathematical or statistical model is obviously preferable (in the macroscopic case) when it is available. There would thus be no real point in trying to validate macroscopic hypothesis within this framework (there may, of course, be some exceptions to this comment) and the mechanistic hypotheses that can be tested are greatly restricted by the form of the environment and the interpretation of events as "motions" in a "plane". We shall indicate the main achievements, so far, and proceed to some radical innovations.

(I) The microscopic behaviour of automata in a potentially cooperative environment has been fully investigated. The initial experiments were rather crude hand simulations, carried out informally by the author. These have been checked and in most respects confirmed by a number of careful hand simulations performed under this contract. The microscopic behaviour of motion, simple interaction, reproduction and the cooperative actions of cooperative groups of automata have been checked on an initial computer programme (designed in accord with the hand simulation) and described in the Annual Summary Report No.1.

- (II) Cooperative groups were shown to replicate as in Annual Summary Report 2.
- (III) Their members were shown to conditionally differentiate. In other words a persistent division of labour has been demonstrated.
- (IV) The stability and cooperative interaction of various types of adaptive automata have been examined. Broadly, an adaptive automaton is able to introduce a form of variation akin to the variation between individuals who have learned from a different background of experience..
- (V) Generative reproduction rules, of the type considered in Administrative Report 5, have been simulated. The later generative simulations deleted type * to avoid the convergence of X and + to a single type.
- (VI) A study of induction with 4 type populations of automata is described in Annual Summary Report 2, and serves to distinguish several mechanisms for induction.
- (VII) Using 4 type populations and a rather simplified environment in which only the main features of the food distribution are preserved (but in which the "payoff" function governing the interaction between the automata is still "superadditive", so that cooperation is still advantageous) it has been possible to consider the effect of constraints that favour the "elaborate" types of automata "x" and "+". These selective constraints are a generalisation of the idea of a "filter" for automata previously mooted in connection with the informal hand simulations. The flow chart for the modified programme embodying the generative rules and the simplified environment - as DIAGRAM 1 - is attached (since it has not appeared in previous reports).

The effect of the generalised "filter" constraint is described in Administrative Report Number 9 and in the present report.

- (VIII) In the present report we consider a predator prey like interaction wherein, by increasing the cost of maintenance parameter value for the "x" and the "+" automata these elaborate types of automata act as predators in respect to the simple automata of types "-" and "1" that act as their prey.

(IX) Implicit mechanisms and explicit mechanisms for population number control or "social homeostasis" have been discussed and exhibited. The "implicit" mechanisms exist in all of the simulations. A macroscopic explicit mechanism has been described and its effects have been observed in conjunction with the "simple" environment programme. This explicit mechanism was shown to stabilise a viable population of automata and to render it "viable" in conditions that proved inadequate in the absence of this mechanism.

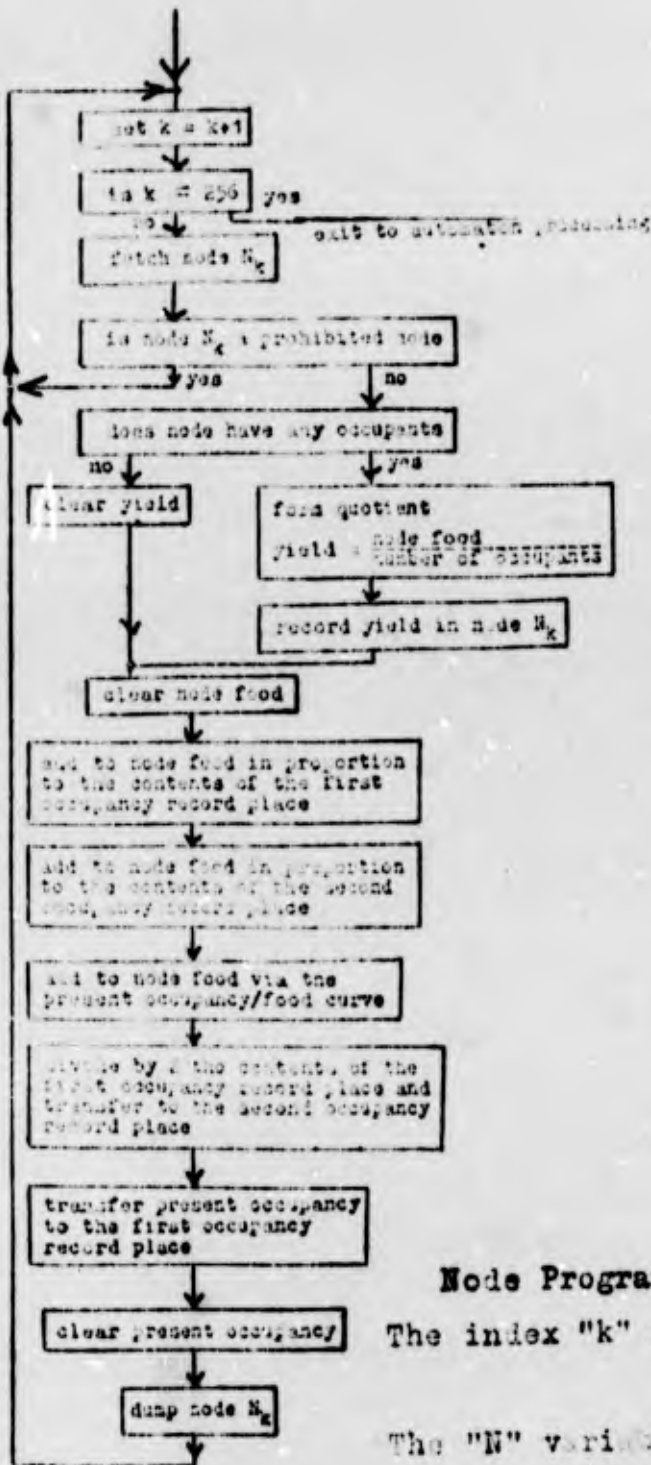
(X) Experiments using a microscopic (and distributed) mechanism for social homeostasis have been performed with a modified type of gregarious automaton. Although the data has recently been extracted from the print out and presented in the present report, the experiments (which involve the original environment) were carried out more than a year ago.

These results have been obtained using a couple of essentially different computer programmes, written by R.J. Feldmann for the I.C.T.1202 computing machine, in machine language. Several modifications of each basic programme have been used. The first programme is outlined in Annual Summary Report Number 1 and replicates the hand simulations that are described in that report. The other programme is outlined in the attached flow chart in DIAGRAM 1. Mr. Feldmann was also responsible for the idea of a gregarious automaton and since leaving this organisation has conducted several "automaton like" simulations (the details of his present work can be obtained from him at the University of Bonn. But, very broadly, I understand that he is trying to apply concepts of the sort used in the present study, within a chiefly biological framework, as a model for particle interactions or at any rate descriptions of particle interactions, in physics).

Although the results we have reported have been presented at various meetings, they have not yet been published (apart from an account of the informal hand simulations and some descriptions of the field). A couple of papers, one chiefly concerned with the induction experiments and the other with the mechanisms for social homeostasis, are being prepared. These

papers may be published in "Kybernetik" or possibly submitted to the "Bulletin of Mathematical Biophysics". A bald statement of the experimental results will appear in the Proceedings of the 4th Congress of the International Association of Cybernetics to complement the description⁵ which has recently appeared in the proceedings of the 3rd Congress, held in 1961.

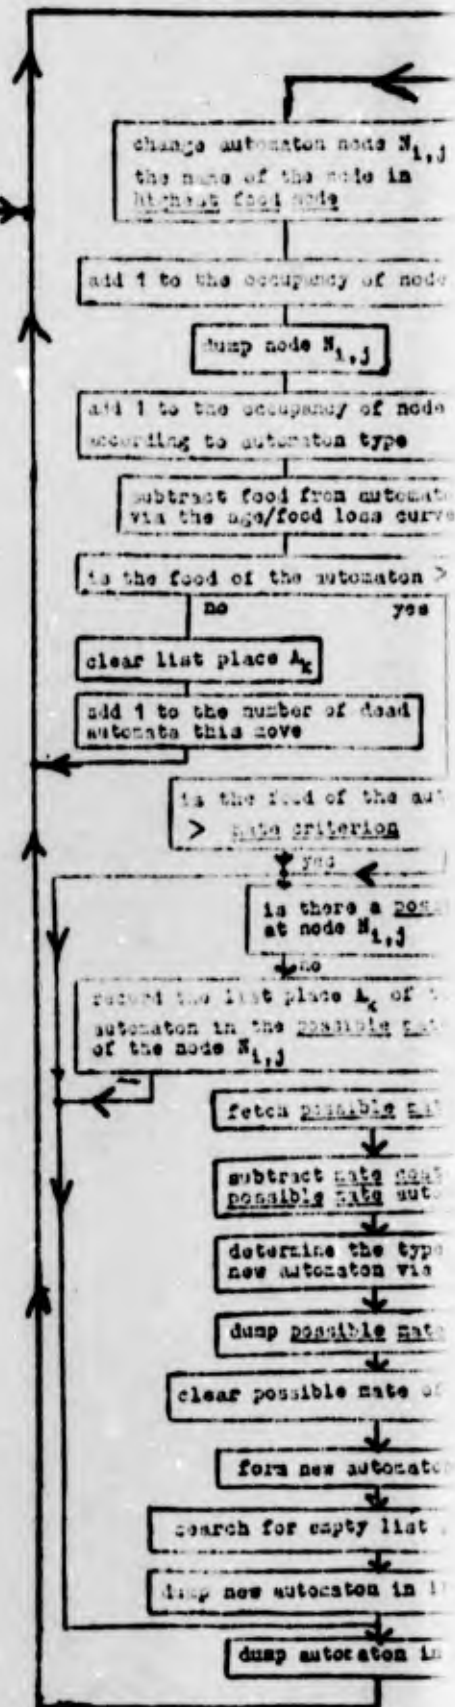
Later, it is intended to examine the already existing print out in greater detail and to publish a commentary upon the more subtle features of the system.



Node Programmes:

The index "k" runs between 0 and 255.

The "N" variable is the node variable




Automaton Programmes:

The index "k" runs between
The "A" variable is the
the automaton resides.
The node variable in the
N_{i,j} i-th row, j-th c


A

GRAPH 1.




GRAPH 2.


NUMBER OF MATCHES, $\circ = -$, $\Delta = +$, $\times = 0$, $\Delta = +$ CENTRALIZED ENVIRONMENT




GRAPH-4.




GRAPH-3.




GRAPH 5.




GRAPH 6.




GRAPH-7.




GRAPH - 8.






GRAPH-9.




GRAPH-10.









GRAPH-13.



GRAPH-14.



GRAPH-15.



GRAPH-16.

Number Automata



GRAPH 17.

Level of food
influx reduced
by approximately
50% at
point -----

Food Available



Number Automata




GRAPH 18

Level of food
influx reduced
by approximately
50% at
point -----

Food Available







Hour 20

GRAPH 19.

Progress of histograms showing the percentage of water that are associated with different values of P_i



Hour 40




Hour 60






Hour 80



Hour 100



GRAPH 20
 Sequence of just-seconds showing the percentage of automation with different values of relative horizontal and vertical motion. The appearance of NCCM Pdk addition to sound is associated with the appearance of another cooperative group of automata.



References:

- (1). Kerner E.H. "Statistical Mechanics of interacting biological species" Bull Math Biophysics 19.121.1957 and Kerner E.H. "Further considerations of statistical mechanics of biological associations" Bull Math Biophysics.21.217.1959.
- (2). Wynne Edwards G. "Animal Dispersion" Oliver and Boyd, 1964.
- (3). Pask G. "A Proposed Evolutionary Model" in Principles of Self Organisation Edited by Von Foerster and Zopf, Pergamon 1961.
- (4). Pask G. "Comments on Evolutionary Systems and self organising Systems" Addendum to paper in Aspects of the theory of Artificial Intelligence, Edited by C. E. Muses, Plenum Press 1961.
- (5). Pask G. "The Cybernetics of Evolutionary Processes and Self organising Systems". Invited address at 3rd Congress of the International Association of Cybernetics. Namur 1961. Published in Proceedings of International Congress on Cybernetics Namur 1965.

Discussion of criteria for designing additional programmes.

In the abstract we aim to simulate a cycle of activity of the form indicated below. It may be variously identified as "learning" or the development of an embryo or the development and evolution of a social structure. This "cycle" is

- (1) Variation, producing a variety of potentially novel forms.
- (2) Differentiation, through a selective process, to produce some stable and distinct entities realising a subset of these forms.
- (3) Dispersion of and discrimination between these stable entities.
- (4) Cooperative interaction between these entities to produce higher levels of organisation, so that variation can provide potentially novel forms of a higher order in an hierarchy of control and organisation.

Depending upon the chosen identification of the Cybernetic Model the variation in (1) may be "Mutation of a genetic code" or "the effect of differing experience"; the differentiation in (2) may be biological differentiation, the application of an available analogy relationship, the realisation of a code, imprinting, or the stable division of labour in society; the dispersion in (3) may amount to adaptation, discrimination, the movement apart of cells from different tissues, or localisation. Finally, the cooperative interaction in (4) may amount to evocation in the embryo of a novel structure by the juxtaposition of previously separate tissues, the production of a novel analogical relationship, or the acquisition of a conceptual structure.

The cycle indicated is entirely consonant with the initial statement of the "evolutionary system" model that is contained in publications such as (3), (4), and (5). The innovations needed in order to bring matters up to date are chiefly a relaxation of the constraints upon the environment in which the

"entities" become dispersed and through which they interact and a separation between the codes for individual entities that develop the "forms" of the previous paragraph) and the codes for the "higher order" entities that (as in the previous paragraph) are produced by their interaction.

We have already commented upon the environmental restrictions. It is proposed to remedy this defect by reducing the environment to a payoff (analogous to food) a communication network (determining the neighbourhood of automata or entities in the simulation) and, a communication cost defined as a function of proximity in this network.

So far as the codes are concerned some further discussion is necessary. In the present simulations we have embedded rules that allow the individual automata to become more elaborate as they develop, for example, that allow pairs of "-" and "l" to form the "x" and "+" types of automata. In this particular case the production of an "x" type or an "+" type of automaton is explicitly an unfolding and realisation of the basic code for an * type of automaton. Hence, the event is distinct from the production of a code of the sort that occurs when a cooperative group of automata develops as a stable system, is replicated, and persists. The cooperative group also has a code although it is not explicitly stated in the programme.

In the first place, we intend to automatically produce or (in degenerate cases) to explicitly embed the codes for higher order structures and to separate the code for the original and automaton like entities (from their realisation) so that this code has the status of a set of instructions that are carried by the automaton. The automatic production of higher order codes presents a chiefly technical problem. It appears to be practicably solvable in several ways but no solution has so far been selected.

The codes of the individual automaton are to be progressively modified as a function of the parentage, age, and experience of the automaton. The parentage determines the potentially available codes. The age leads to a decrease in the "competence" (in the biological sense of this term) of the automaton (the number of codes that can be realised is decreased with increasing age). The experience of the automaton or within the biological nomenclature, the "field" effect of neighbouring automata, acts upon the codes that are competent at the particular instant when an experienced

event takes place. This leads to a more elaborate individual automaton than those used in the present simulation and as a practical issue to populations of fewer automata. So far as the elaboration is concerned the form of an automaton is a subroutine in the programme that is elaborate enough to specify the learning process associated with a single subskill of a structured skill and the first flow chart for this individual, shown in DIAGRAM 2, has been designed with this identification in mind. Hence, the increased elaboration has the effect of allowing for something comprehensive enough to embody the previous algorithmic learning model as the code for an individual automaton (the adjunction of the mechanistic model to produce a realised entity or automaton is exhibited, in a slightly degenerate form, by DIAGRAM 2).

Obviously the (necessarily reduced) numbers of automata will present some difficulties. It seems likely, however, that we can compensate for an inadequate size of population by introducing statistical variables that represent the rather loosely coupled effect of an aggregate of the other members of an imaginary large automaton population. Once again there are some technical problems and apart from the case in which the simulation is identified with the learning process in a student who is acquiring the subskills of a structured skill these problems have not been solved. The solution used in this case may be universally applicable. It is presented in the paper "Man as a system that needs to learn" that is cited in Part 4 of this report and is being embodied in the programme of which the flow chart above represents a single subroutine. This programme is being written for the 1301 computing machine at Brighton and, in a specialised form, in Sweden also.

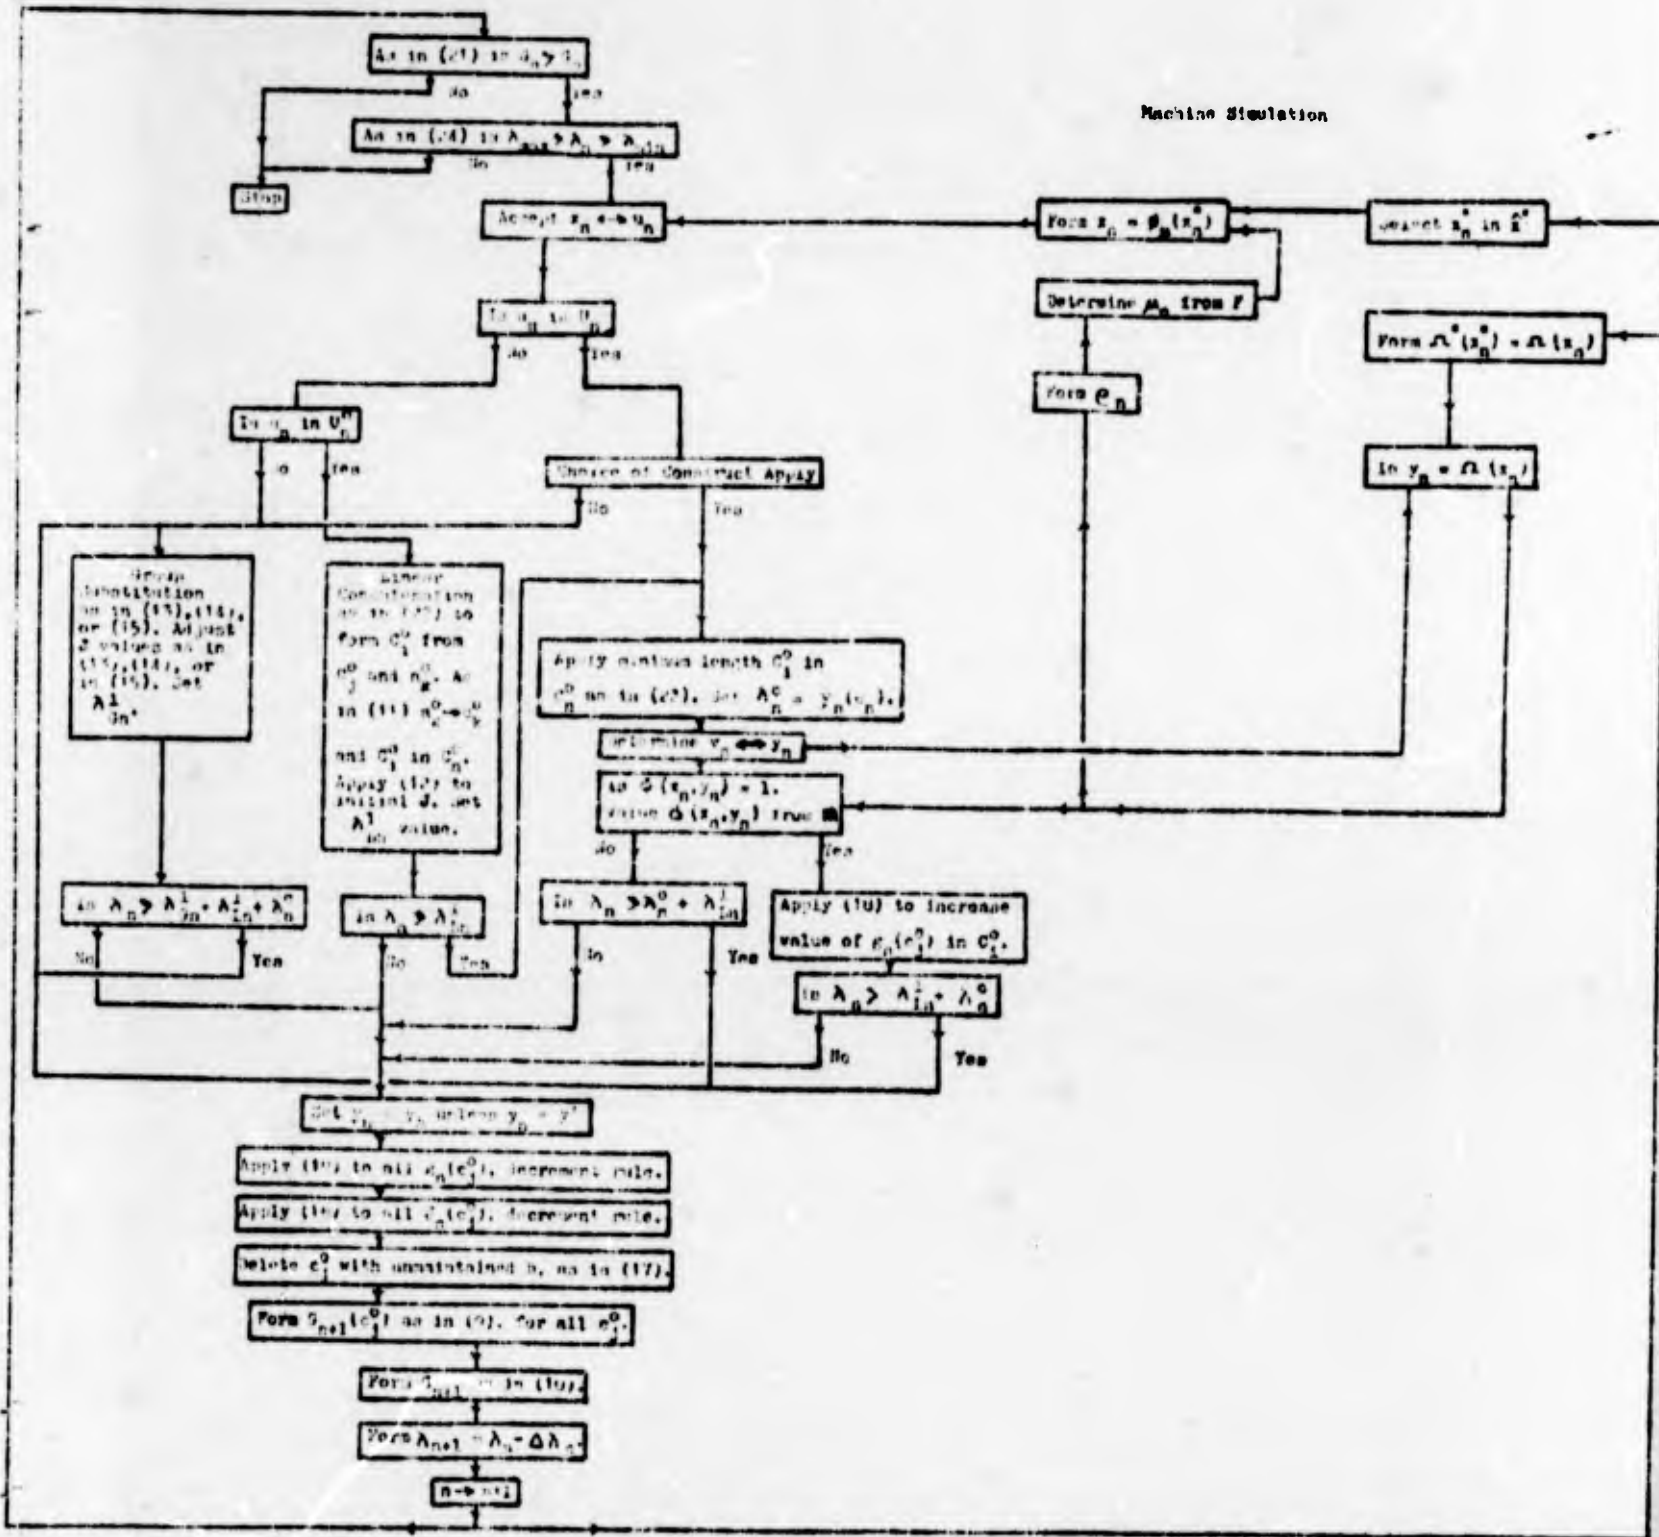


DIAGRAM 2.

SYMBOLS USED IN THE FLOW DIAGRAM OF THE
COMPUTER PROGRAMME.

- U_n - A set of intelligible problems.
- u_n - An intelligible problem.
- U_n^u - A set of understandable problems.
- u_n^u - An understandable problem.
- U^* - A set of similar and unsimplified problems.
- u^* - An unsimplified problem.
- x_n - A stimulus (which may be simplified).
- x_n^* - An unsimplified stimulus.
- \hat{x}^* - An arbitrary sequence of x^* .
- v_n - A solution .
- v^+ - A correct solution.
- y_n - A response state.
- y^+ - A correct response state.
- η - A solution mapping from u_n to v^+ .
- η^* - A solution mapping from u^* to v^+ .
- ϕ_{A_i} - A simplification procedure of degree A_i .
- $\delta(x_n, y_n)$ - Knowledge of results signal; $(x_n, y_n) = 1$ if $y = y^+$ and 0 if not.
- ρ - Proficiency.
- a_k^o - A code unit.
- A^o - A string of code units.
- c_i^o - An embodied code unit.
- C_n^o - A string of embodied code units.

Symbols for flow diagram (cont.).

- b° - Minimal subsystem of mentation.
 λ_n° - Number of operations used in L° application.
 λ_{Ln}^1 - Number of operations used in Linear Concatenation (L^1 applications).
 λ_{Gn}^1 - Number of operations used in Group Substitutions (L^1 applications).
 $\lambda_{max.}$ - Maximum number of performable operations in time Δt .
 $\lambda_{min.}$ - Minimum number of performable operations in time Δt .
 $\Delta\lambda$ - Change in $\lambda_{max.}$ and $\lambda_{min.}$.
 $\varepsilon_n(c_j^{\circ})$ - "Activation" of c_j° .
 $J_n(c_j^{\circ})$ - "Concentration" of c_j° .
 $G_n(c_j^{\circ})$ - $\varepsilon_n(c_j^{\circ}) - J_n(c_j^{\circ})$.
 G_n - $\sum_j G_n(c_j^{\circ})$.
 G_0 - Limiting value of G_n .

The subscript 'n' refers to the nth. trial,

Part 4.Further Development of the Cybernetic Learning Model.

Since Annual Summary Report Number 1, the concepts underlying the basic learning model have been considerably extended and attempts have been made to identify the abstract Cybernetic structure with physical systems. So far as its psychological identification is concerned the argument is set out in "Man as a System that needs to learn" (this paper has been modified to bring it up to date, it is, at last, genuinely in press, and a Xerox copy of the final manuscript has been provided for the file). A partly psychological and partly "data processing" identification, also up to date, is available in "Some Comments on the organisation of Men, Machines, and Concepts" which appeared in "Education for Information Science", Edited by Heilprin, Markussen, and Goodman, Spartan Press and MacMillan 1965 (reprints of this paper are also in the file).

Finally, I have attempted to produce a form of this model devoid of ordinary symbolism; in other words, with mathematical symbolism replaced by a graphical convention. This graphical form is adequate for several biological or ethological identifications of the model and, in particular, for an account of the conjoined maturation and learning process called "imprinting". This development of the model appears in a book that is being written (by the present author) called "The Cybernetics of Living Systems". Since some innovations are involved and since the "imprinting" model seems novel enough to deserve simulation (it is certainly amenable to simulation) Chapter II and Chapter IV of "The Cybernetics of Living Systems", which contain the relevant material are attached to the present report as appendices. Although the content of Chapter II and Chapter IV is unlikely to be changed, these appendices should not be regarded as the final manuscript. In several places the style needs to be improved before publication, but the basic ideas are presented.

APPENDIX I

CHAPTER II

Cybernetic Models and Control Systems.

Control Systems.

Chemical plants and factories are controlled by automatic devices which are designed to maintain the goal of satisfactory operation. Aeroplanes are flown by automatic pilots that respond to haphazard perturbations of the vehicle and satisfy the goal of keeping it on a prescribed course. The temperature in a building is often adjusted by a control mechanism that senses the room temperature, compares it with a standard value and turns the heating on or off as an operation that achieves the goal of maintaining the room temperature value nearly equal to the standard.

Nature is teeming with control systems. Large scale control systems maintain the character of the Biomes (such as "Coniferous Forest" or "Desert" or "Grassland") given the basic resources of air, sunlight, water, minerals and space in which to grow. The action of these ecological control systems can be observed in the microcosm of an aquarium that is freshly inoculated with pond water. After several days, the aquarium attains a dynamic equilibrium condition. The plants, animals and bacteria live in balanced numbers; the animals use the oxygen produced by the plants, (which aerates the water) one species uses another for food and waste products, in non-toxic concentrations, are reutilised and transformed. The whole thing is driven, through photosynthesis, by sunlight. We can show that it is a control system by artificially disturbing the balance when, within limits, the balance is automatically rectified or by changing the basic conditions of illumination or size when a different dynamic equilibrium is achieved. Hence, the entire system is stable for it achieves the goal of maintaining a balanced population in the presence of arbitrary disturbances. Further, it is ultrastable, (in the sense of Ashby) because, when the basic conditions are modified, it achieves another goal which is compatible with membership in the class of all viable aquaria.

Distributed Control Systems.

Unlike the majority of mechanical control systems, an aquarium is a distributed control system. It would be difficult to identify and localise the part which senses a deviation from the goal condition or the part that operates (like the heating apparatus in room temperature control) to approximate the goal. The difficulties over localisation are exacerbated by the complexity of the system; it is a mesh of control rather than a single control system. There are the metabolic control systems of the individual organisms, the control systems that maintain the concentration of various products in the water by changing the secretion and ingestion of the population as a whole, and many others all of which interact with one another. But it will often be impossible to put a finger on the components of a single control system in this mesh, even if it is isolated. The functions of different parts are carried out upon different occasions by different physical elements. Further, on any one occasion, a functionally distinct action or computation may be carried out in many different places. Hence, we say that the control system is distributed (or that the functionally distinct entities making it up are distributed).

Distribution of control is neither confined to natural systems nor is it true that all natural control systems are of this form. There is, for example, one class of error resistant automata that is functionally distributed. These automata are networks with the property of automatically compensating for arbitrary changes in structure (that, in the absence of a special design, would introduce error into the overall computation). These automata have been considered in detail by Cowan and Winograd; their design entails the recoding of an original network which is able to carry out a computation C without error (using components that each do a specific job) into a further network in which C is computed by components that may do several jobs. Some mechanical and some natural systems are of this sort, for example, some parts of a mammalian brain.

Some are distributed without the special virtue of error resilience. Other natural systems, including many of the control systems in a brain or in animal behaviour, are perfectly localised.

Basic Organisation.

Whether or not its parts are localised, the organisation of a control system is well defined. It is built up from the functional entities in DIAGRAM 8. namely,

(1) A bounded domain representing properties of the environment or of the system itself, upon which the system operates and in respect to which it aims for a goal.

(2) A sensory or recognition device able to determine states of this domain, a goal specification, and a comparator that determines the value of a "distance" or "deviation" from the goal (the extent to which the present condition of the domain is unsatisfactory and must be modified).

(3) A source of variation. In the simplest control systems (such as temperature controllers and positional servomechanisms) the variation occurs in the domain; the room temperature fluctuates or the position of the motor shaft changes haphazardly. In other cases there is an inbuilt source of variation which, in the absence of any change in the domain, perturbs it by making trials or performing experiments. We shall call these active control systems and some systems, for reasons we shall discuss later, are necessarily active. The locus of the variation becomes a matter of definition when the system is distributed; the source of variation needed to keep an aquarium going may be in the environment or in the control system as we please, for the domain is not well localised.

(4) A decision rule that stipulates, for any deviation from the goal, what sort of operation shall be performed in order to approximate the goal. The decision rule usually belongs to the familiar class of "negative feedback" rules; the system annuls deviation from the goal by operating to a degree that is the inverse of the deviation or "feedback" signal.

(5) An operation that acts upon the domain to achieve and preserve the goal state. The diversity of operations is as great as the diversity of domains; turning on a heating apparatus, changing the concentration of a metabolite, moving a muscle and, for an interesting class of systems that have a collection of individual organisms as their domain, changing the rate of growth of these organisms or their rate of reproduction. The crucial points are (I) that the operation shall be throttled back by the instruction received from the decision rule (in the absence of which the muscle may move a large load or the reproduction rate may be gigantic) and (II) that only an energetically or numerically small change of instruction from the decision rule is sufficient to produce an energetically or numerically large change in the operation. Properties (I) and (II) yield the essential condition that there is an amplification in the control loop from sensor to effector or, in a slightly esoteric usage of the term a power gain.

Artificial and Biological Control Systems.

Probably the most cogent distinction between man-made and natural control systems is that, in the former case, the goal is specified by a designer, whereas in the latter, it is discovered by an experimenter. Neat as it sounds, the distinction is no more than a useful generalisation; some man-made control systems (such as the evolutionary simulations in Chapter VI, and various "learning" and "intelligent" artifacts) do not have any well defined goal; conversely, the experimenter often uses a design that constrains the natural system and partially determines its goal.

Although the goal of a control system at a given instant is fixed, it need not represent a static condition. It may, for example, stipulate a constancy of rate, (in which case the state of the system is not constant, by definition), or it may prescribe a restriction upon the changes that occur without in any way, deterring change; the dynamic equilibrium of an aquarium is an illustrative case; the maintenance of a constant number of individuals in a population wherein some members decay is another.

Stability.

The control system is stable if it can achieve its goal. Now stability is an important but deceptively simple concept. Strictly it is defined with respect to a given class of disturbances or perturbations, with respect to a particular state description (a model in which certain properties of the system and certain operations are deemed relevant and used as descriptive coordinates) and often with respect to a particular class of laws (such as the class for which all changes in state can be described by linear equations). Control engineers can usually make enough assumptions to justify fairly strong stability statements and there is a collection of more or less stringent "stability criteria" which allow them to inspect certain properties of a system and to determine without setting it in motion, whether or not it will be stable. Most of these "stability criteria" are inapplicable in biology because the basic assumptions cannot be justified. Some of them can be applied to special cases (like the control of voluntary movement and the Mantid control system in Chapter III, and those (such as the Lyapunov criterion) which do not depend upon linearity assumptions can be used more generally.

In the absence of formal criteria it is still possible to say something about the capabilities of a stable control system. It must be able to compensate for any admissible deviation from the goal; it must be able to act more rapidly than the environmental fluctuations that perturb its domain) i. must be able to operate in an appropriate fashion. The most comprehensive formulation is given

by Ashby in his law of "Requisite Variety". If we know the possible operations of the control system (any one of which may constitute a long sequence or strategy of actions, it is unambiguously stated) these can be used to label the columns of a table; knowing the possible perturbations of the domain, these can label the rows of this table; the entries (or row and column intersections) are possible outcomes. The goal of the system is satisfied by a collection of one or more of these outcomes. Ashby points out that a control system is only able to achieve its goal and thus to be stable if the possible variety of the actions (or columns) is at least as great as the possible variety of the perturbations (or rows) in this tabulation and this is its requisite variety.

Composition Rules for Control Systems.

The control system in DIAGRAM 8, is the basic unit in a Cybernetic Model and it can be combined to yield the reducible structures shown in DIAGRAM 9 and DIAGRAM 10. Since these structures are reducible they are also Cybernetic Models.

DIAGRAM 9, is a parallel combination of systems a, b, and c. These systems are coupled in a specific fashion that does not entail the modification of their goals, and they interact. Most of the control systems in an aquarium are coupled like this. Alternatively, a, b, and c, could represent separate automata communicating by signs.

DIAGRAM 10, is a sequential concatenation of control systems or an operation sequence (of the operations carried out by the control systems). As Brissey points out, this structure is characteristic of the motor and "conative" functions of an organism (whilst the parallel combination is characteristic of sensory and cognitive functions. We shall also come across operation sequences in many behavioural organisations.

At the motor level the control systems A, B, and C, might concern lowering and taking body weight on the right foot, raising the left foot and bringing it forward. The operation sequence (A, B, C) images a crude sort of walking movement. In any operation sequence (walking or not) we call the goals of A, B, C, subgoals and engage B when A's subgoal has been achieved

and C when B's subgoal has been achieved. An engagement of A thus realises the walking movement (providing that the contingent subgoals are achieved) which is the goal of the entire operation sequence. Now A or B or C in isolation are largely irrelevant to walking; C is mechanically impossible without B; and the engagement of A, B, and C, in an improper order is usually absurd. Thus the improper ordering A, C, B, will land an organism flat on the floor.

Codes for Control Systems.

If we have an alphabet of signs for objects and operations then this organisation of (A, B, C) which expresses the contingent dependence of the several control systems and leads to the satisfaction of the "walking movement" goal, may be stated as a code expression such as "Achieve Subgoal A, go to B and achieve subgoal B, go to C and achieve subgoal C". This code does not stipulate how the subgoals should be achieved or what physical operators realise the operations prescribed by the code. It specifies, in each case, some operation and this in turn, a class of operators. Hence the code (or the class of codes we consider) is the explicit statement of a goal (and the implicit statement of its constituent subgoals) in a given sign system or language. The code, in fact, is a programme like a computer programme, wherein the constituent codes for A, B, and C are subroutines. In deference to the fact that codes such as the code for A, B, C represent control systems, Miller, Gallanter, and Pribram have called them TOTE or "Test, operate, Test, exit" units (the "exit" specifying the address for a further subroutine of the same sort). The operation sequence is a physical realisation of a TOTE unit sequence (each operation being replaced by a specific operator).

The embodiment of a code prescribes something or describes something; it can accept instructions, such as "realise the operations that are prescribed, as operators", it can also issue instructions, such as "consider the next code sequence"; hence, the embodiment of a code, its written form, must have certain properties. There must be an alphabet of signs able to denote objects or other signs and there must be rules for the association of signs and the denotation of associated groups of signs. Finally, there must be an interpretation mechanism whereby sign sequences are converted into actions or further signs.

Various physically stable structures serve as signs and are used to embody codes, for example, chemical molecules and parts of the nervous system.

The Control of Control Systems.

The construction of DIAGRAM 11 (I), is an hierarchy involving the control of control. It is obtained by replacing the domain of any control system (the higher level or L^1 system) by the collection of goals available to an L^0 or lower level control system (so that the L^1 controller acts to change the goal of an L^0 controller and receives feedback information regarding the achievement of these goals).

If, as in DIAGRAM 11 (II) the L^1 control system is embodied in a physical system but the L^0 control system is not, then the L^1 control system exerts certain constraints upon the L^0 control systems that may subsequently be embodied; in particular, it determines their possible goals. However, no L^0 system exists. In deference to this fact, we adopt the symbolism in DIAGRAM 11 (II) and comment that the constraints exerted by the L^1 system upon the putative L^0 system are identically an L^0 code for this control system. It is entirely possible, if the system is realised in a fabric with malleable or inertial characteristics, that these constraints persist after the L^1 control system has been active or even after it has decayed. In this case we adopt the symbolism in DIAGRAM 11 (III) which is equivalently an assertion that the code for an L^0 control system is written in an alphabet of L^0 signs.

The control of control is alternatively represented in terms of selection. The necessary construction is obtained by replacing the L^1 operation of L^0 goal variation by an L^1 operation of selecting an L^0 control system (able to aim for only a single goal) from a collection of alternative L^1 control systems and it appears in DIAGRAM 12 (I). As indicated by DIAGRAM 12 (II) the selective action of an L^1 control system determines an L^0 code and if the L^1 control system of DIAGRAM 12 (II) decays (although its selective constraints persist in a stable fashion) this is an L^0 code written in an L^0 sign system as suggested by DIAGRAM 12 (III).

The selective construction for representing the control of control, is particularly useful in connection with the genesis of operation sequences. Amongst the most common L^0 operation sequence codes are programmes, produced by an L^1 control system, that specify the order in which certain operations must be performed without determining in detail what operations shall be performed. Such a code is indicated by DIAGRAM 13 (I); it is identical with the specification of Z and the domain and ordering of the boxes labelled as A, B, and C; the exact contents of A, B, C, the operations to be performed, are not yet determined because the specifications of the subgoals "A", "B", and C only determines a class of control systems. Hence A may be legally substituted by A_1 or A_2 , B by B_1 or B_2 and C only by C_0 (hence, C is determined, which is exceptional) to yield the operation sequence in DIAGRAM 13 (II). The legality of any substitution is merely a matter of whether or not it satisfies the L^0 code and any other constraints that may be imposed.

At this stage (A, B, and C have not been substituted) an instruction to achieve the goal Z is ineffective. But it may (and commonly it does) provide an input to a further L^1 control system, M, that is able to select substituent operations within the L^0 code constraints often also imposing constraints of its own. The system M need not be the L^1 system that erected the initial L^0 operation sequence code and usually it is not. Further, in some important cases, the substitution process depends upon externally provided L^1 signals and upon L^0 signals that are dealt with by any L^0 systems that have already been constructed as a result of partial substitution. (In any case, as Brisseley also points out, the construction of a sequential process entails some 'parallel' data).

The genesis of an operation sequence is readily generalised to yield parallel combinations of control systems and thus amounts to a general model for learning in the non-trivial sense of building up control systems. To obtain the learning model we have invoked a malleable or inertial fabric, (or the concept of adaptability) but the control hierarchy is reducible to control systems (embedded in such a fabric) and consequently this model for learning is a reducible Cybernetic Model.

To complete the picture we comment that in experimental situations the L^0 signals which act upon the learning process are stimuli and that the L^1 signals are "reinforcements" or "rewards" (where applicable, instructions also).

If the adaptable fabric has the physical characteristics required for a sign system (according to our previous account) then the codes we have considered may be written or embodied in this fabric. Given such an embodiment the concept of a level of control or of an organisation, the hierarchy L^0, L^1, \dots , is open to an equivalent linguistic definition. These levels of organisation are distinct domains of expressions in the sign system; L^1 is the domain of expressions that describe or prescribe L^0 expressions and L^0 is the domain of expressions that denote reality hence, if the sign system is a language, a level of control is simply a level of discourse in this language.

Systems that are and are not immutable structures.

Any localised machine (mathematically, any finite automaton) must decay due to haphazard abrasion of its structure. So far as man made machines are concerned, we accept and neglect this fact. Until they wear out, their fabric as well as their organisation is regarded as immutable. The same sort of decay affects the localised living organism. Its fabric is obviously subject to the inroads of age and of predation. But the actual rate of breakdown of its constituents is very high indeed. An organism could never achieve its normal age (indeed it could never be born) without special arrangements for reconstructing the physical embodiment of its organisation. The components of an organism are in continual flux. Even in the higher animals the materials forming many of the organs are replaced completely every few weeks and, in some cases, every few days.

To exemplify the permanence of an organisation within a labile fabric we shall consider the control systems of a cell, where the flux of material is rapid and obvious. Although the organisation of a cell is complex it has recently become rather well known because of several important and lucidly described advances in molecular biology. A study of this complex organisation also provides an especially straightforward illustration of the difference between (and the reality of) levels of organisation and control such as L^1 and L^0 ; it will be evident that these are levels of coding (for the L^1 and L^0 alphabets of signs and the forms of the corresponding expressions are unambiguously defined); and we shall thus demonstrate the plausibility of a later contention that L^1 and L^0 are languages, or alternatively that they are levels of discourse in a rather comprehensive single language.

The Cellular Control System.

The lowest level control systems in the cell are organised enzymes* or sequences of enzymes. The organisation of these catalysts may or may not entail spatial distribution on a membrane (it does, for example, in the case of the respiratory chain mechanism in the mitochondria). In any case, the essential organisation is a chemical specificity wherein the product of a catalysed reaction is used as a feedback signal to inhibit at least one of the enzymes catalysing this reaction. It was once believed that inhibition depended chiefly upon substrate competition (in other words, that the product molecules competed with the substrate for the active catalytic site on the enzyme). Some inhibitory systems do work in this way, for example, Oxalacetate ions inhibit succinic dehydrogenase in the Krebs cycle of DIAGRAM 15. On the other hand, many enzyme molecules are "allosteric". They have more than one "active" site. One site is

*Biological catalysts made of protein.

occupied by the substrate molecules, the other is occupied by a specific molecule that alters the configuration and activity of the enzyme. It may act as an activator or as an inhibitor, for example, L Isoleucine inhibits the action of L Threonine Deaminase which is an enzyme rather a long way back in the reaction chain that gives rise to L Isoleucine. The important point about allosteric interactions is that the several specific inhibitors and activators are signals in a communication system and need bear no immediate relation to the substrate or the product of the reaction that is modified. Hence, cooperative effects are able to form parallel and sequential combinations of control systems and these cooperative interactions bind the metabolic control systems into a coherent organisation (this level of organisation will be denoted as L^0 , and the metabolic control systems will be called L^0 control systems).

Enzymes are proteins (molecules principally made up from long sequences of amino-acids). There are 20 alternative amino-acids and since their pattern determines a protein they form an alphabet of 20 letters in which the L^0 code for a protein is written as a long sequence. Now proteins decay and must be replaced by a manufacturing process. Enzymes and other proteins are manufactured at structures in the cell called ribosomes, which also participate in the process of decoding genetically specified information that gives instructions for assembling amino-acids into protein molecules.

The system concerned is shown in DIAGRAM 14. The genetic code which specifies the cellular organisation is embodied in a substance, D.N.A. which makes up the chromosomes in the nucleus of a cell. The D.N.A. is made from chains of fairly small molecules called nucleotide bases (thymine, adenine, cytosine and guanine, labelled as T.A.C.G), that are held together by sugar molecules and are associated with proteins called "histones". At this level of organisation, which we denote as L^1 , the instruction for taking up and assembling a single amino-acid appears as a triple selected from T.A.C and G. so that a complete protein is represented by a string of triples. Any D.N.A. strand specifies many proteins, which are thus specified in an L^1 code as words in a four letter alphabet. We comment that D.N.A. strands exist (in the cells of nearly all organisms) in pairs and that the pairing is orderly; if strand I contains C at a given position, Strand II contains G, if I contains G then II contains C, if I contains T then II contains A and if I contains A then II contains T; hence the strands are complementary.

Now the decoding system which interprets L^1 instructions in terms of the L^0 specification of enzyme proteins involves the fact that one D.N.A. strand can form a complementary strand of a substance called "Messenger" R.N.A. Like D.N.A. the strand of R.N.A. is made up from four bases; A.C.G. as in

D.N.A. and uracil U, in place of T, but in R.N.A. the bases are held together by a different sort of sugar.

The production of an informationally homologous and complimentary strand of messenger R.N.A. is effected by a polymerising enzyme, but it is controlled by a number of distinct L^1 control systems. In a multicellular organism the cells are differentiated; only some of the genetic code can be copied. They are also restrained by hormones which, possibly acting through the histone proteins, allow or disallow the production of R.N.A. from specific loci on the D.N.A. molecule. In addition there are L^1 control systems that are able to recognise L^0 signals. The underlying L^1 mechanism consists in the "structural" loci or "genes" on the D.N.A. that are able to produce messenger R.N.A. codes for a particular system of enzymes and another pair of loci, the "operator" gene that control the structural gene group and determines whether or not it will produce the messenger R.N.A. that it is able to produce and the "repressor" gene that manufactures a specific material the regulator, which acts (1) upon the operator gene, which it inhibits or disinhibits according to whether it is or is not activated, and (2) as an L^0 signal recognising device in the sense that it is specifically activated by some molecule that is either one substrate or one product of the enzyme system*. If the specific activating molecule activates the repressor and inhibits the operator to "switch off" enzyme production, the control system is called repressible; if the converse effect "switches on" enzyme production it is called inductible; in each case the effect may persist after the L^0 signal has appeared and is no longer present. (The underlying mechanism, at any rate for repressible systems, is often less elaborate insofar as the operator gene and the repressor gene are combined at a single locus).

The messenger R.N.A. that bears the L^1 code for an enzyme group is broken down fairly rapidly in the cell, chiefly by an enzyme, ribonuclease. Before this occurs the messenger R.N.A. passes on its L^1 code to a ribosome which translates it and interprets it as an L^0 code and which embodies this L^0 code in the protein molecules of the enzyme.

*Footnote. The inductible system that has been most thoroughly described is responsible for producing the enzymes, Permease and Galactosidase which are required for the assimilation and metabolism of lactose in the bacterium E Coli. This E Coli lactose system is induced by the L^0 signal "lactose" in the culture medium. The most thoroughly described repressible system manufactures the enzymes needed to synthesise Tryptophane and it is repressed by the specific L^0 signal of "Tryptophane" in the cellular environment.

The ribosomes are made from protein and a stable form of R.N.A. ribosomal R.N.A. (which is also manufactured from a D.N.A. code but which bears no message). In the cell there is another stable form of R.N.A. (again produced by D.N.A. but bearing no L^1 message) which is called transfer R.N.A. The transfer R.N.A. exists in at least 20 specific forms corresponding to the 20 different amino-acids that are used for the manufacture of proteins and that embody the signs in the L^0 alphabet. In the cell, each of the amino-acids is attached by a specific enzyme to a specific transfer R.N.A. through an energy bearing phosphate linkage. The Transfer R.N.A. amino-acid combinations are assembled on a ribosome to form the L^0 coded strands of enzyme proteins that are specified in the L^1 code by the messenger R.N.A.; the ribosome thus "translates" the L^1 code into an L^0 code, or it "interprets" this code; further, it cooperates with an energetic mechanism to embody the L^0 code as a specific protein.

The control of this interpretation process is mediated at the point where informationally homologous messenger R.N.A. is produced from a D.N.A. coded genetic locus. This "between level" control of an L^1 process by an L^0 process has already been described (in terms of inductible and repressible systems) and given a specific signal molecule for each L^0 control system, it couples the L^1 control systems to the L^0 control systems to form an hierarchical organisation as in DIAGRAM 14.

Reproductive Process.

Amongst other materials the metabolism of the cell produces nucleotide bases and the remaining constituents of D.N.A. Although the exact mechanism is not yet known, it is evident that a signal is produced whenever either (1) The cooperative interaction between the parts of the cell is impaired by decay or (2) the multicellular organism needs more cells of the type concerned; in the latter case the signal is hormonal. (In passing, it is useful to compare (1) with the recycling signal in the Tsombega system of Chapter VII and (2) with the signals in Chapter VIII). On receipt of this signal the double strands of D.N.A. are replicated using the materials provided by the metabolism and the genetic code is thus freshly embodied in a copy of the original D.N.A. After D.N.A. copying the entire cell divides and the copied D.N.A. conveys the genetic data to the daughter cell.

In principle the copy may be perfect; an organism acquires potential immortality in a static world insofar as it has the trick of using its own metabolism to reconstruct its own code in a stable embodiment like D.N.A. and thus to reproduce. However, such an organism could not keep pace with a dynamic world; to do so, it must be capable of variations (that are supplied in fact, by genetic mutation) and it must be equipped with a "natural selection" principle whereby those variants that fit the changes in this world preferentially survive.

Energetic Metabolism.

The whole process depends upon a supply of energy which, in animal cells, is mainly derived from the oxidation (or equivalently the dehydrogenation) of foods such as sugars and fats. Now the synthetic activities we have outlined (as well as the more specialised cellular activities of actively maintaining an ionic balance or, in muscle cells, of contracting) accept energy in a common form and unit. The energy is accepted from a common energy carrier, Adenosine Triphosphate (or A.T.P.) which denotes an energy rich phosphate linkage as the common unit and is thereby converted into Adenosine Diphosphate or A.D.P.*(1) The basic energetic requirement is thus to manufacture A.T.P. as the "charged" form of energy carrier given A.D.P. as the "discharged" carrier that has activated some reaction. To do so entails the provision of a highly structured physical system that oxidised food in stages that provide the right size of energy units. Most of the parts of this system situated in the mitochondria of the cell as indicated in DIAGRAM 15*(2). It thus appears legitimate to conceive the primary function of the L^0 control systems as the construction and maintenance of this energetic machinery and the synthetic machinery (such as the ribosomal mechanisms supporting membranes and so on) which it feeds.

Hence we arrive at a reciprocal correspondence that images the crux of hierarchically organised control in systems that (unlike most man made artifacts) are not immutable and must actively maintain their characteristic patterns. The L^1 control systems construct and maintain the embodiment of the L^0 control systems; the L^0 control systems replicate, freshly embody and reproduce the L^1 control systems. Finally, in an environment that changes and calls for an adaptable collection of L^0 control systems, the variation of this environment must be countered by a variation amongst the L^1 control system codes. This variation is provided by the "mutation of genes" whereby the genetic code is modified in a manner that is not directly correlated with changes in the environment.

*Footnote (1). This is a method for activating a chemical reaction that is catalysed by some enzyme, there are some exceptions to this statement, for example, some reactions accept from A.D.P. replacing it by Adenosine Monophosphate.

**Footnote (2). Since the mechanism of energy transfer is peripheral in our discussion, no description of the process is attempted. However, DIAGRAM 15 offers a fairly up to date picture of the mitochondrial organisation which will be intelligible to anyone versed in the nomenclature of this field. At least it serves to indicate the highly organised and physically structured pattern of the L^0 control systems that are maintained against continual abrasion and decay.

The Statistical Mechanics of the Cellular Control Process.

Any cell contains a vast number of control systems of the type we have described, any one of which is conveniently represented in the notation of DIAGRAM 16. The problem of describing their conjoint activity has recently been tackled by B.C. Goodwin and his theoretical analysis is one of the outstanding achievements in Biochemical Cybernetics (it gives rise to a "statistical mechanics" of cellular interactions). We shall not deal with the mathematical part of the model but will briefly indicate how the model has been assembled since the analytic process is complimentary to the linguistic approach we have adopted (in common with the majority of authors who have modelled cellular reactions, we have taken an L^1, L^0 , distinction for granted. Similarly Waddington, for example, makes a distinction on empirical grounds, between the "epigenetic" level of control and the "metabolic" level of control. Goodwin arrives at this hierarchical structure in a different way.)

The state description of the cell contains variables that represent the concentrations of different reactants, a pair of variables being needed to describe the state of each control unit. Under the influence of specific catalysis and given the energy required, these concentration values are altered to generate a trajectory of states of the cellular system (its state at any one instant being represented as a point in a space of the descriptive variables). Now, on further scrutiny, it turns out that this state description can be partitioned. This possibility arises because the systems controlling the reactions have relaxation times that are not continuously distributed (roughly a relaxation time is the interval needed to regain equilibrium conditions after a unit disturbance). Some systems have relaxation times which, in the cells of higher organisms, are in the order of several hours; others have relaxation times in the order of a few minutes, and there are few in between values. It happens that the systems with long relaxation times often correspond to "epigenetic" systems concerned with the synthesis of large molecules at an L^1 level of control; similarly the "metabolic" systems at the L^0 level of control often have short relaxation times. But even in the absence of this correlation, it would still be possible to partition the variables in the state description for, although these different types of systems interact, the variables of the long relaxation time systems appear as fixed parameters in the equations describing the activity of a short relaxation time system and conversely, in dealing with the equation for a long relaxation time system, it is possible to assume that all of the short relaxation time systems that are associated with it can reach dynamic equilibrium between one change of its state and another. The distinction of levels in DIAGRAM 16, is thus based upon the observation that the relaxation times of the cellular systems are grouped rather than being continuously distributed and it is possible

to separate the essential variables of systems having different relaxation times as subspaces of the space of descriptive variables and to order these subspaces under logical inclusion. Thus the subspace of the long relaxation time (as it happens, epigenetic) system variables include the subspace of short relaxation time (as it happens, metabolic) system variables. But, since the epigenetic variables are parameters in the metabolic equations and the metabolic variables are essentially stationary valued in the epigenetic system equations the analysis of the epigenetic systems is not necessarily more elaborate than that of the metabolic systems.

The control units in DIAGRAM 16, are plausible entities just because of this separability. It is evident that any control unit is specified epigenetically but that it interacts in a not very well determined fashion with a host of other control units through the metabolic processes it controls. Since these interactions are at the metabolic level we may regard them as weak interactions to be specified, ultimately, in no more than a statistical manner.

In addition, there are strong couplings between the control units which, as suggested in DIAGRAM 16, are mediated by repressor feedbacks to specific genetic loci. These interactions must be stated explicitly and they will be viewed as deterministic in character and will appear in the equations of activity of the control units.

The equations of activity for any single control unit can be deduced from some plausible assumptions about the biochemical processes involved; so can the equations for certain strongly coupled combinations of control units (some combinations are excluded because the equations cannot be integrated; hence the exclusion is a matter of mathematical convenience and it could be avoided by using and possibly developing more elaborate mathematical procedures). Goodwin shows that any control unit acts as a restrained oscillator and that strongly interacting combinations yield elaborate oscillatory behaviour with long term temporal organisation.

The difficulties that beset an analysis of the behaviour of the entire cellular system are (I) that although all of the control units can be described within the same state description their initial states are undetermined and (II) that the weak interactions between the control units are not well defined.

To avoid the problems engendered by (I) and (II) Goodwin has

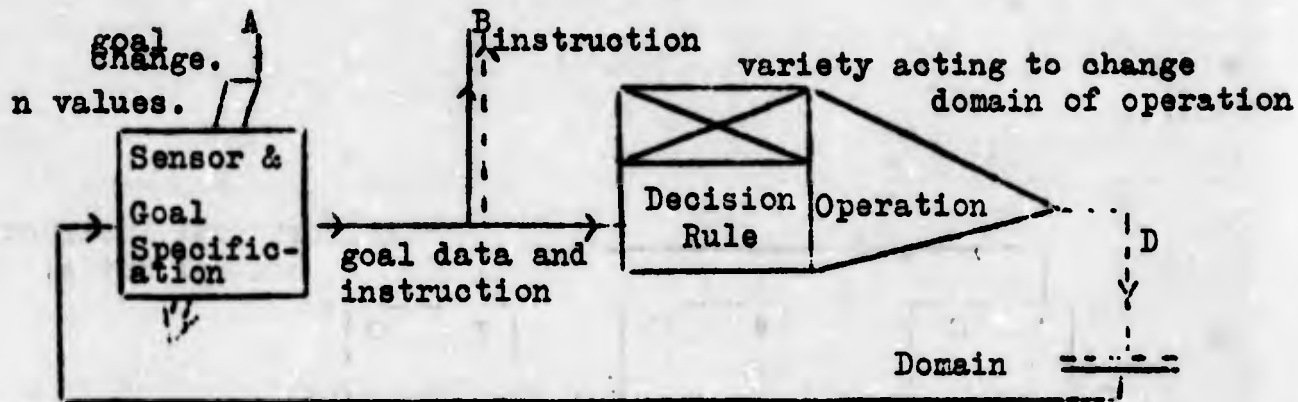
introduced a statistical mechanics* as in Chapter I. In place of a single control unit we consider an ensemble of control units which may either be conceived as an indefinitely large collection of abstract replicas as in Chapter I, or as the large but finite collection of physical systems in the cell. The states of the ensemble are represented as points in the descriptive space and the trajectories of the members of the ensemble are associated with a constant, in time, of this motion, which (by analogy with the statistical mechanics of gas molecules, rather than control units) has the characteristics of energy. To avoid confusion between the energy of molecules and the constant of the trajectories of the ensemble of control units this quantity is called the "Talandic Energy" of the ensemble. Although the "Talandic Energy" is a constant for the entire ensemble, it is the commodity which is exchanged, when there is coupling, between subsystems of the ensemble.

Another important statistical parameter is a measure of the excitation of the entire system or of a subsystem of control units which (again by analogy with the statistical mechanics of gas molecules) is a "temperature". This "Talandic Temperature" determines the expected direction of exchange of "Talandic Energy". It is also possible to show that the "Talandic Temperature" represents a measure of the degree of non-linearity of the oscillatory behaviour of the control units and to argue that the maximum temporal organisation of such a system increases with its "Talandic Temperature".



Footnote.

The prerequisites are (1) Statement of the differential equations for single and strongly interacting combinations of control units and their presentation in Hamiltonian form (2) Identification of a suitable integral of the equations for an ensemble of control units with the energy (3) Statement that the equations for an ensemble satisfy Liouville's theorem so that the density of neighbour points about the state point of any representative member of the ensemble is invariant (4) Derivation of class of averages over the ensemble equivalent to time averages and assumption that all admissible initial states are equiprobable.


BLANK PAGE




ABREVIATES TO;



And if goal change instruction is "aim for" or "remain inactive" only, this abbreviates to;



or to:



Notice that n of the II units are equivalent to one n goal alternative I unit. Thus for n = 5,






DIAGRAM 8. I, II,




PARALLEL COMBINATION OF
CONTROL SYSTEMS.

DIAGRAM 9.




Operation sequence A,B,C, with subgoals A,B, and C, and a goal Z. Notice, (1) that domain of control units may be the same as in A and B, or different as in the case of C, and (2) that the operation sequence is the operation of control unit Z.

DIAGRAM 10.



The construction for control of control and hierarchical organisation. If L¹ control system is embodied in a physical structure, but the L⁰ system is not embodied, then the constraints exerted upon subsequent realisation or activity of an L⁰ system, are an L⁰ code for its L⁰ goals. The existence of a code is represented by the symbolism:-



When as in learning, the L¹ system operates upon a malleable or plastic fabric, the constraints may persist as a "written" code, even if the L¹ "writing" system decays. In this case we use symbolism:-






DIAGRAM 11. I, II, III.

I.



An alternative "control of control" symbolism in which the L^1 system selects amongst alternative L^0 control systems. If only one L^0 system is selected from a set of alternatives as determined by the L^1 source of variation, this set, due to the L^1 constraints, is an L^0 code which is conveniently represented as:-

II.




If in learning, the L^1 system imposes its constraints by "writing" the L^0 code in some malleable or plastic fabric, then the "written" code persists after the L^1 system decays, as in :-

III.



DIAGRAM 12. I, II, III.

As a special case, the code constraints imposed by a higher order control system, may determine the order of operations in an operation sequence leaving more or less latitude regarding the operations, as indicated below. We comment that Z cannot achieve it's goal until the operation sequence code is substituted by real operations. The substitution process is carried out by M.



A = A₁ or A₂ B = B₁ or B₂ C = C₀

When the operations of this programme or operation sequence code are physically realised as operators, the form of the operations is determined. The actual substitution depends upon an L¹ control system, M, that selects from the available set, and upon the L⁰ interaction. In either case the result is an operation sequence whereby Z can achieve it's goal.






DIAGRAM 13. I, II.

Diagrams 14 and 15 are too complicated to reproduce accurately on stencils. Therefore, they have been omitted from this report.


Single unit
Goodwins symbols



Strongly coupled pair
of units in Goodwins symbols



Single control unit
in present symbolism



Strongly coupled pair of control units
in present symbolism




DIAGRAM 16.

APPENDIX II

BLANK PAGE

CHAPTER IV

Fabric and Organisation of Memory.

The Malleability of Animal Fabrics.

Although animals perform some wonderful feats of learning and memory, it is prudent to reserve the wonderment for the organisation of the process. There is nothing very astonishing about the fact that animals retain an impression or "trace" or "engram" of their history for they are built from a malleable fabric which is inherently capable of habituation and adaptation, and as Ashby and Beer have pointed out, so is the fabric of many natural systems. We are so accustomed to machines, such as computers, wherein the components have been designed to avoid the effects of habituation and adaptation that the ubiquity of these processes in nature is often overlooked. In fact, it is difficult to imagine conditions in which a neurone, for example, would not bear many persisting traces of its previous stimulation and the real problem is to select a relevant class of mechanisms of learning from an unduly rich collection of possibilities.

Mechanisms of Malleability.

Four sorts of mechanisms have been widely discussed as possible vehicles for the malleability in the nervous systems of the higher animals. These mechanisms chiefly involve the neurones of the brain but there is an impressive body of data suggesting that the glial cells which support and nourish the neurones also take part in memory processes.

(1) If a loop of self excitable neurones is once stimulated it will continue to reverberate and the existence of this activity signifies its previous stimulation. A suitable network of self excitable circuits may be thrown into different stable modes of activity by distinct stimuli and the pattern characterising its prevalent mode signifies the previous stimulation. Dynamic memory systems of this type are "cleared" of all the retained information if their electrical activity is interrupted and because of this it is known that long term memory is not dynamic though short term memory (and the mechanism for placing data into long term memory) almost certainly is. Thus brain freezing (which brings the electrical activity of the brain to a standstill) has no effect upon at least some long term memory, nor does the

administration of a disruptive electrical shock. But either freezing or shock impairs the assimilation of experience over an interval extending up to about half an hour before the treatment and for some time after it.

(2) If a neurone, i , is excited by signals arriving at its j -th synapse it is plausible, on the grounds of physical chemistry, to suppose that a change in membrane properties takes place which signifies, to the i -th neurone, the fact that it was once excited by a signal at the j -th synapse. An older and more extreme version of this hypothesis suggested that activity dependent growth of nerve fibres established or extended synaptic connections. Whilst there is plenty of circumstantial evidence in favour of at any rate the modern version of this hypothesis, it is difficult to conceive an experiment that would test or disprove it.

(3) The members of another class of hypothesis, suppose that excitation of a neurone by a temporal or spatial pattern of impulses P , gives rise to a physio-chemical change in the cell membrane that is peculiar to P , rather than the fibres that conveyed P . Once again, the hypothesis is almost certainly true but it is difficult to disprove.

(4) A peculiarly stable type of P sensitive memory has been proposed in connection with the experimental work of Hyden and others. To be definite we present a mechanism for memory which goes considerably beyond the experimentally verified field. Suppose that a spatial or temporal pattern of impulses P induces the production of a specific messenger R.N.A. which is used to synthesise a particular protein, p , with the property that in contact with the pattern P , protein p breaks down into substances that excite the neurone. Once that p has been produced it remains available because a messenger R.N.A. is available for its synthesis. Finally, we assume that several messages, P_1, P_2, \dots, P give rise to several specific proteins p_1, p_2, \dots, p and that a given neurone develops a particular class of proteins able to "recognise" a corresponding class of patterns.

The chief evidence for this sort of memory mechanism comes from experiments in which rats learned to perform special skills, such as

balancing on a tight wire. In rats the learning of a balancing skill depends upon cells in Deiter's nucleus and it was found that nucleotide base ratios in the R.N.A. from Deiter's nucleus cells of rats that had learned to balance on a tight wire differed very greatly from the normal base ratio (or the ratio for rats that had been simply suspended and rocked), hence suggesting the learning controlled synthesis of an abnormal R.N.A. Similar results have been obtained for other animals, but they are open to the criticism that the base ratio analytic technique demands separable cells and the acquisition of rather special experimental skills. Again, learning can be inhibited, in certain conditions, if an animal is treated with an antibiotic such as Puramycin, that inhibits protein synthesis but these results can be criticised because, apart from learning, the treatment interferes with other aspects of the metabolism.

It seems likely that all of the mechanisms (1), (2), (3), and (4) are at work in normal learning. Data must be consolidated in short term memory by some dynamic process. An appreciable delay is involved in producing an R.N.A. molecule and although reverberation patterns could hold the information whilst the data in short term memory is assimilated into long term memory it is likely that an intermediate "plastic" change is involved. Surprisingly small "plastic" changes of connectivity are needed to stabilise a particular mode of autonomous activity in a network such as DIAGRAM 25, in the sense that once the i -th reverberation R_i has induced certain changes of connectivity C_i , the most probable pattern of reverberation becomes R_i . In other words, if the network characterised by C_i is subsequently excited in a haphazard fashion it will prove to reverberate in the i -th mode R_i . Further, a system of this sort admits a type of contingency which is needed in any learning system. The changes in connectivity required to achieve R_i in DIAGRAM 25, are only admitted if R_i , arises from a goal approximating form of excitation of the network.

The Location of Memory.

People ask "Where is the memory?" meaning, "Where is some item stored?" and they ask "Where is memory?" meaning, "Where is the memory organ?". It may be the case that neither enquiry is answerable, but in each case something can be said. The original picture of memory, as a sort of file index system with something like a cortical neurone registering a single item was reversed by experimental work, such as Lashley's that showed learning and memory in rats and other creatures to be unspecifically effected by large cortical ablations. The file index theories were replaced by dynamic memory theories in which items were registered by complex modes of activity. These notions have been rejected so far as long term memory is concerned although they plausibly apply to short term memory.

The present thinking on this subject is that long term memory is neither localised nor is it diffuse in the dynamic sense. To illustrate the point, consider a "phantom limb" (somebody with a painful and diseased arm has it amputated. Later, he feels the pain that used to exist, say in his fingers, even though the hand is missing. If so, he is said to have a "phantom limb" (which is a curiously tangible sort of memory). It is possible to remove the phantom limb by destroying a very precisely definable tract of nerve fibres that are responsible for signalling the pain (destruction of the tract can be achieved by a beam of ultrasonic sound or in other ways). The "phantom limb" disappears. Hence, the memory must be localised. But a few months later, the limb returns. Hence the original memory was evidently represented in some coded form, in another part of the nervous system and, in this sense, it was diffuse. The relearned "phantom limb" can again be exercised by an operation that is entirely local. But a few months later it is relearned, so the relearned memory was also represented in some distant part of the nervous system.

Briefly, a memory has a specific operational locus, which is one of its representations. But it has many, and conceivably an indefinitely large number, of representations in the brain. There

is a great deal of evidence in favour of the view that memory involves a continual relearning in which whatever is relearned becomes differently and more diffusely represented. We comment that an organisation of this type is identical with an associative memory system in which data is stored in terms of many tokens for its properties. The datum "Aunt Amelia's sitting room" is stored, for example, in terms of properties, such as being rectangular, being a rectangular room, containing an aspidistra plant and belonging to an irascible old lady. For each property value there will be a token or tag and "Aunt Amelia's sitting room" is ultimately represented as a way of associating its tokens. As more items are stored in an associative memory, more descriptive properties are generated and insofar as these are relevant to an existing item, this item becomes differently represented due to the acquisition of novel tokens that relate it to the novel descriptive properties.

The Items in Memory.

What "item" is remembered and relearned. The evidence points very strongly to the proposition that an organism remembers what to do about a situation in order to achieve a goal; in other words, the basic "item" is a control system code which, depending upon the level of abstraction, (or the level of the coded control system in an hierarchy of control) is either an instruction for doing something goal directed or a prescription for building an operator that will do something (or a prescription for a class of prescriptions and so on). Consequently, the descriptive properties in a memory system should be either properties of goal directed action or more abstracted properties of properties of action, and there are good reasons for supposing that they are.

Aunt Amelia's aspidistra is recalled as that into which, we might pour a glass of tonic wine and Aunt Amelia's irascible disposition is recalled as that with which we argue to get our own way. The case of the phantom limb appears to contradict this hypothesis (it seems as though a phantom limb is a "fact" or simply a "feeling"). But even here it is feasible to argue that the feeling occurs because the organism has a memory that makes it try to do something or to make a compensation that cannot, in fact, be done or made

(which avoids the rather unfair repartee that a phantom limb is a pathological memory anyhow).

In support of the contention is the fact that the most basic sorts of learning, including coordination skills and maintenance of posture, are impaired or prevented if the subject is not allowed to do what he normally would do or if the motor feedback derived from these actions is distorted or interrupted. On a slightly different tack, all learning which has been examined experimentally is associated with goal achievement. In laboratory experiments of the type we shall consider in Chapter V, the animal is rewarded by the experimenter for doing whatever it should learn to do and there is little doubt that the creature learns how to achieve the goal of getting a reward. The experimenter must, of course, find out what the animal treats as a reward; this may be something obvious, like food for a hungry animal or it may be something that satisfies a more subtle physiological requirement; indeed, it may be any commodity that the animal has a "drive" to get. From the tenor of our previous discussion it will be evident that not all "drives" are of the directly physiological type (subserving a reduction in hunger or the provision of sexual satisfaction). If the creature contains or acquires any code or programme which prescribes a goal, then there is a perfectly good symbolic drive to achieve this goal (when the programme or code is acquired as a result of previous learning the reward suited to a symbolic drive is often called a "secondary reinforcer", in contrast to a "primary reinforcer" which acts as a reward for a physiological drive). Symbolic drives, at any rate those that are innate, can often be generalised. Most of the higher animals have a generalised curiosity drive, and Desmond Morris has emphasised the "neophylic" or "novelty loving" propensity of certain species. None of this is at all surprising. We might predict that animals would seek to explore their environment on the grounds that they are active control systems; we might expect to find higher animals exploring in the hinterland between the unfamiliar and the positively frightening. But these symbolic drives, however predictable, are not physiological drives in the classical sense. I am stressing the point because all of the counter-examples to the present contention namely all the exemplars of learning that is not

goal directed rest upon the absence of physiological drive and rely upon the supposed absence of any drive. The chief exemplars are latent learning, copying, and imprinting. Of these latent learning and copying almost certainly represent delayed forms of imprinting which is a process we shall consider in a moment.

Where is the Process?

Is the organ of memory the brain (as we supposed above) or something more than the brain or some part of the brain? If we accept the view that memory involves relearning, then there is no real answer to this enquiry; certain sorts of relearning depend upon the states of the organism and others upon states of its environment (yet it would hardly be reasonable to call the environment part of the memory organ). It is perfectly reasonable even so, to say that some cells in the organism are specialised as malleable units. According to most present day thinking, these are all neurones, and the neurones specialised as malleable units are scattered rather diffusely about the nervous system. In some animals these cells have been identified, for example, J.Z. Young has advanced a scheme of the sort shown in DIAGRAM 26, to explain how the octopus learns to advance or retreat from a visual form. The specialised cells in DIAGRAM 26, are the amacrine cells in the octopus optic lobes.

Not surprisingly, if the "items" in a memory are "codes" for control systems, the organisation of memory appears to be identical with the organisation of a goal directed hierarchical control system. At the lowest level of organisation a control system realised in physically malleable fabric could adapt in favour of its goal, so that if some environmental condition is repeated it will subsequently select an action that is goal approximating. Thereby it acquires or improves its decision rule. If we start with an active control system as in DIAGRAM 8, then this adaptation becomes a primitive sort of learning. The source of variation guarantees that

if the system is not perturbed by its environment then it will act upon its environment (in a fashion that is haphazardly selected within the constraints imposed by its structure). Such a device is imbued with the rudimentary auxiliary goal of learning (in a rather inefficient, trial making, way) about its surroundings.

The "control of control" construction yields a higher level of organisation wherein the L^1 control systems learn about and operate upon (according to DIAGRAM 11) or select and construct (according to DIAGRAM 12) the L^0 control systems in their domain. If the attitude controller of DIAGRAM 21, is embodied in a malleable fabric it "learns attitudes" and is, of course, a special case of DIAGRAM 11. If each parameter value is associated with a distinct L^0 control system, then DIAGRAM 21, has the form of DIAGRAM 12. Since all of these structures are reducible to active control systems of the sort in DIAGRAM 8, the learning process is also reducible.

When embodied in a malleable fabric "control of control" is a non trivial sort of learning; an hierarchical concatenation of trivially "learning about learning". Since the process entails the stable association of control systems, or the codes for them, it is apt to call it "concept building"; notably "concept building" in a semantic universe of discourse is isomorphic with the reciprocal process of cellular reproduction, when the system is interpreted in Chapter II, in a biochemical domain.

If "control of control" is iterated, the resulting system is able to build an indefinitely large hierarchy of concepts. Each level in this hierarchy is concerned with a distinct collection of signs and may be regarded not only as a distinct level of discourse but as a distinct sign system or internal language. Hence, the interpretation of concepts, codes, or one level of expressions in terms of another level (or in terms of the system language) entails translation. If, as in DIAGRAM 27, a pair of organisms A and B are assumed to communicate concepts at level L_j then

they must each be equipped with an interpretative mechanism that translates between L^j and their system language.

Several Levels of Discourse.

There are a couple of outstanding and closely related questions about learning and memory. One question is "To what extent can a pair of organisms A and B communicate concepts?". This is really an issue of what interpretative codes A and B jointly possess. The other question is "To what extent does learning depend upon the codes that are innately determined (specified as part of the genetic instructions of an organism)?". This will be dealt with later.

If A and B are to communicate as suggested in DIAGRAM 27, then they must be equipped with arrangements for translation from the higher conceptual levels into the system language. Since the system language certainly accomodates L^0 signs for data about the environment, this may amount to providing an L^0 interpretation or translation of expressions at any other level of discourse (we could achieve L^1 communication by establishing a direct connection between the L^1 organisation in A's head and the L^1 organisation in B's head, using the L^1 data codes. But this would imply "getting into" A's head and "getting into" B's head in a rather literal fashion which is physically impracticable).

*Footnote. The distinction of levels in this rigid way is not accidental. It stems from the need, in a reducible model, to avoid any ambiguity regarding the level of discourse to which statements refer. A similar distinction is made, in a logical context, between "logical types" and yields a "type hierarchy" or its equivalent.

Do organisms need to communicate higher level concepts? It appears that they do. The teaching relation "A teaches B a skill" certainly entails L^1 discourse, for A must be aware of the goal to be satisfied by B, he must extensively define this goal in L^0 , provide exemplars of the stimuli that are dealt with in the skill (again L^0) and he must reinforce B if his responses are goal satisfying. The relation "A teaches B a skill" occurs very commonly in animal communities and is about the simplest case of higher level discourse. The offspring of ravens, for example, grow up in an elaborate social milieu which evidently, from Lorenz's account, involves a great deal of teaching on the part of the parents and other adults.

Do organisms have the necessary translation systems? They do, insofar as they are able to indulge in any higher level discourse. But there is plenty of evidence that the acquisition of these translation systems (usually by a teaching process) is something of a stumbling block and we shall argue that it imposes the chief restriction upon social evolution. In the absence of a translation mechanism a concept is isolated and cannot be shared.

Is there any way, apart from direct connection (which is impracticable) of establishing higher level discourse between A and B. The reply is "No", if stratification is retained and if, as a result of this, statements in the system language are unambiguous. The reply is "Yes", if the stratification is avoided by using a self referential system language that is able to describe and construct its own expressions. In particular, it is possible to specify the entire concept building process in a suitable self referential language and the stratification becomes unnecessary. The price to be paid is an irreducible ambiguity of reference. Thus the expedient of discarding stratification is only open to organisms able to tolerate ambiguity or to give their own idiosyncratic interpretation to statements in the system language.

The consonance of Codes.

Let us now return to the question of which codes are innate. Any animal has some behaviour that

is innate, in the sense that just this behaviour would develop even in an odd and atypical environment. Most animals have some behaviour that is learned in the sense that an adaptation occurs, which is jointly consonant with its innate codes and the structural codes of its system language. Some of birdsong is innate (the recognition of certain sound coordinates and the use of certain muscular actions) but some of it is learned (the recognition of specific sound patterns and the refinements of singing). The same comment applies to most of sexual behaviour and to the selection of a habitat in which to live. The really notable point is that the behaviour of the animal ultimately fits the structural code of the system language and thus that there is a "consonance" between the genetic code of the animal and the linguistic code. In fact, the "consonance of codes" is a general biological principle and the sequential learning process that brings it about is of general importance. The consonance between the innate codes of the whole animal and the structural code of its language is due to a particular manifestation of this sequential learning process called "imprinting" which occurs in the first part of life.

Some Experiments on Imprinting.

Imprinting was first noticed by Spalding in 1872. The phenomenon was studied by several workers including Lorenz in the late 1920's and it has been intensively investigated by Tinbergen, Thorpe and others, so that a large body of data is available. Ideally, and as originally conceived by Lorenz the phenomenon of imprinting consists in the rapid and irreversible acquisition (without the satisfaction of any physiological drive) of an attachment or an aversion usually entailing the recognition of a releaser sign, which is, by definition, in the system language. Imprinting of a given type can only occur within a short "sensitive period" in the animal's life.

Typically the sign that is recognised is "Mother". In normal conditions the baby animal comes across its real mother within the "sensitive period" for maternal imprinting and, if so, it subsequently exhibits maternally directed behaviours like "following". But the sign for a mother is not innate, although the

need to have a "Mother" sign is innate. Hence, the imprinting process can be "fooled". Many animals can be imprinted to "follow" and to espouse rather inappropriate objects, for example, ducklings that are exposed to the experimenter in their sensitive period will subsequently follow the experimenter; Harlow's baby monkeys became maternally attached to a cloth covered wire structure; by an unfortunate accident a peacock was once led to regard a turtle as its only possible mate (imprinting applies to sexual as well as maternal attachments). In fact, it applies to most signs and actions that are important for survival or that are used in connection with other animals of the same species).

In a recent survey of the field, Sluckin has pointed out that imprinting does not usually have this ideal character but merges into ordinary learning by trial and reward. It is not so immediate as used to be imagined and there is some evidence that it may not be any more irreversible than other learning. So far as drives are concerned, the curiosity drive is presumably satisfied if it exists. But in any case, the goals are perfectly adequate symbolic goals that are achieved by the fact of imprinting, as such. Finally, the work of Guiten shows that in chicks, at any rate, the sensitive period is more adequately conceived as an ordering of sensitivities. The organism is sensitive to Type A imprinting until A₁ is imprinted and to type B imprinting until B₁ is imprinted.

Perhaps Sluckin's main point is that imprinting occurs within a framework (of perceptual attributes and predetermined action patterns). Chicks, for example, become imprinted to follow (as they would their mother) after any moving dark or light object, or even any flashing light. This, and a few other features of a mother are salient properties to which they attend and these perceptual attributes are inherited. They are also, of course, consonant with the system language indeed they determine its universe of discourse. Whilst accepting this point of view, we comment that the number of perceptual attributes normally associated with the imprinted figure may be very much larger than the number encountered in the laboratory. In some respects, the relatively simple figure of Harlow's cloth covered wire structure did not act as a mother;

later in life the baby monkeys exhibited several deficiencies. By hypothesis some of the relevant perceptual attributes were missing from the maternal figure in the experiment.

The Sequential Development of Consonance.

The consonance of codes developed through imprinting involves a sequential control process which belongs to a rather large class of similar processes, all of which can be developed from the paradigm case in DIAGRAM 13. As a first step in this development we convert the substitution process, M, in DIAGRAM 13 into an operation sequence. This yields the organisation of DIAGRAM 28 (I).

The best exemplar of this organisation is the metamorphosis of insects; such as the familiar transformation of an insect egg into the larval form of caterpillar and the caterpillar into the adult which is a butterfly or a moth. Within hours after the insect egg begins to divide an initial embryonic control system, R say, lays down the code, for the adult form; this is embodied in scattered groups of cells called imaginal buds. As the larva develops (enlarging chiefly by undifferentiating expansion of the existing cells and accommodating this expansion by moulting), the imaginal buds remain dormant. Now R also produces a substitution system, M, that is embodied in the hormonal control mechanism in DIAGRAM 29. The secretory cells in the brain produce hormones of which one stimulates the secretion of a further hormone P.G.H. from the insects prothoracic gland, (as in DIAGRAM 29. there is also an inhibitory feedback of P.G.H. to the brain. Similarly, a pair of bodies, the corpora allata that are closely associated with the brain (and are analagous in insects, to the pituitary gland) secrete a juvenile hormone J.H. In the presence of J.H. In the presence of J.H. the P.G.H. system controls the cyclic changes of moulting, and the imaginal buds remain inhibited. As the concentration of the corpora allata's J.H. decreases* the pupating insect larva becomes responsive to external stimuli such as temperature which render the brain and the

*Footnote. Although the mechanism for the decrease in J.H. concentration is not elucidated the process can be reversed experimentally, for example, the insect can be retained as a giant developing larva by injecting J.H. or made to metamorphose prematurely by removing its corpora allata.

P.G.H. system functionally equivalent to M. A sequential process occurs in which the imaginal buds are activated and the code α is substituted stage by stage. Coincidentally, the larval cells are dissolved and used to provide nourishment for the α coded adult.

The next modification of the paradigm organisation appears in DIAGRAM 28 (II). Now the initial process R lays down a pair of codes α and β as well as M. Although the substituents for the α and β operations are undetermined before M becomes active, R does determine the interconnection of the domain of the putative control system.

The best exemplar of this organisation is the biological specificity of motor and sensory limb innervation patterns which has been chiefly studied in the embryos of amphibia. α is a central representation of action patterns (such as the walking movement) or of sensory patterns. β is a peripheral representation of these patterns which names the muscles and nerves involved (and is probably a chemical coding). Rather late in embryonic development M, which is the limb innervation system, sequentially substitutes the operations in α and β in a consonant fashion, and brings about the required neural connections. The character of this process is revealed by the fact that it can be experimentally "foiled". By ingenious surgery that alters the peripheral "names" it is possible to induce a bizarre α to β correspondences reflected in anomalous structures such as amphibia that have futile movement patterns or that scratch their backs when they are tickled on their front.

Finally, consider DIAGRAM 30. The codes α and β and the substitution system M, are again laid down by R, which is however, conceived as an evolutionary rather than a genetic control system. β is the set of constraints upon a system language, its syntax and alphabet. α is a code for releasers and other signs in an individual organism with the system language β (so that, in a sense β could be specified by an aggregation of the α codes characterising all members of some local population). The substitution system M is the imprinting system. Although M only acts upon α we stipulate in DIAGRAM 30 that it receives at least some external L^1 control signals (in other

words, at least some β signs can be interpreted or recognised by the organism as instructions to M). These signals have the form "thus L^0 input belongs to the class of mothers" which may amount to no more than "the class of stimuli associated with a flashing light". For imprinting, nothing further is required. The imprinting of a given element of α is restricted by the admissible β signs in (unless, experimentally, we import some bizarre object with the same name in the β universe of discourse; say, the experimenter himself). The process is sequential (B can only be imprinted after A has been imprinted) the goals are internal and symbolic goals; but notice that the organisation would depict reinforcement learning rather than imprinting if we had required an effective L^1 feedback of data. Indeed, there is a sense in which the specificity obtainable by reinforcement feedback is replaced, for imprinting by the restriction of β .

To illustrate the point, DIAGRAM 31, is an imaginary but plausible substitution plan. It is assumed that α admits $A = A_1$ or A_2 or A_3 or A_4 , that $B = B_1$ or B_2 or B_3 or B_4 and $C = C_1$ or C_2 . The code β imposes the partially sequential dependencies shown in DIAGRAM 31, which indicates the constraints imposed upon the imprinting process by different experience of β signs emitted by other organisms or the environment.

The Production of Novel Organisations.

So far, we have presupposed the existence of higher levels of organisation without specifying how they appear. The omission is deliberate, for an adequate account of the development of higher levels of organisation entails the cooperative interactions which will be discussed in Chapter VI. However, this is a logical point to consider the essential part played by differentiation in providing the entities that do cooperate presupposing some vague idea of what cooperation is, and leaving it vague until Chapter VI). The imprinting process is one of the most important types of differentiation (it might, alternatively, be dubbed "symbolic" differentiation). A more elaborate but also more tangible differentiation occurs amongst the cells in a developing embryo, leading to the production of coherent tissues like skin and brain and kidney. Indeed cellular differentiation and imprinting

are isomorphic if we accept a view, most explicitly stated by Flickinger, that the development of "competence" in a cell is a sequential process. Although rather little is known about the detailed mechanism of cellular differentiation it is commonly agreed that cells differentiate because of an interaction between their "competence" and the embryonic "field" in which they reside. The competence of a cell (at a given instant) is the collection of genetic codes or genetic loci that can be realised at that instant. Initially, the cell is equipotential (in the sense that any of its codes can be realised). It becomes restricted, Flickinger suggests in a sequential fashion, as certain genetic codes are blocked or prohibited. If α in DIAGRAM 30 represents the genetic code then M is analagous to this sequential process.

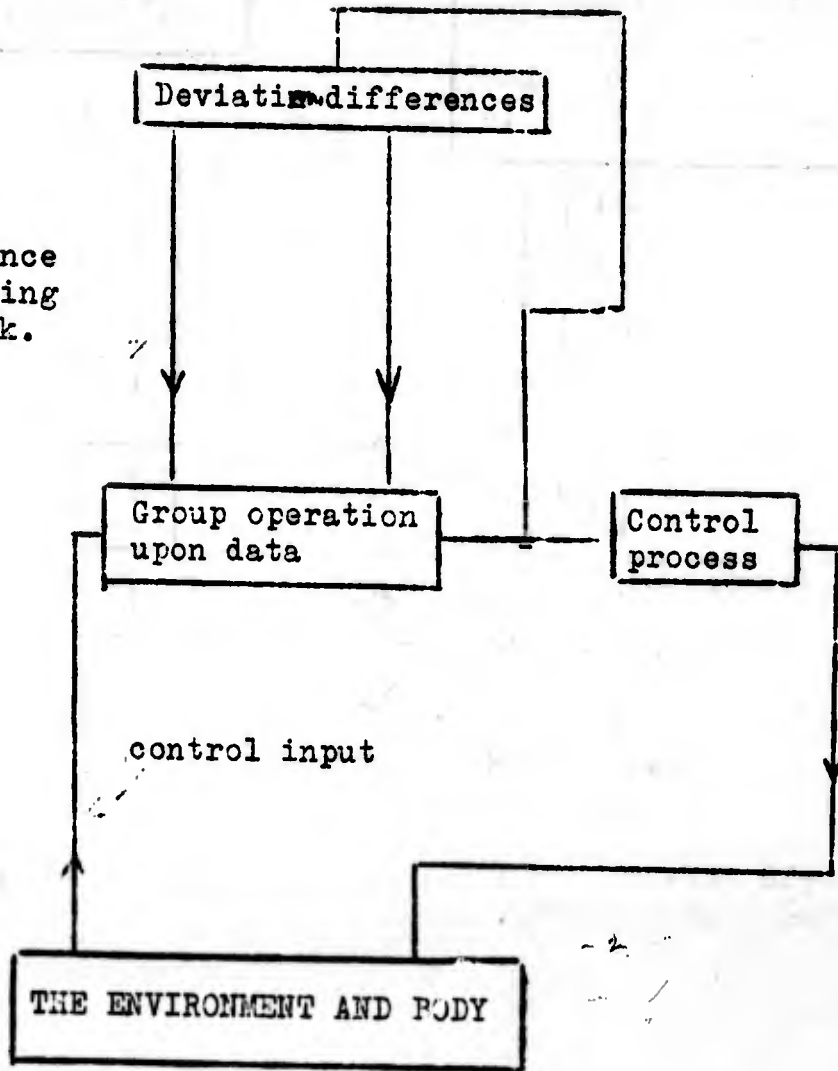
The embryonic field of a cell consists in the influence exerted upon it by all neighbouring cells in their immediate condition of competence. The L^0 discourse in DIAGRAM 30 is analagous to the β restricted influence of other members of the local population upon the individual organism, whilst the L^1 discourse is strictly analagous to the effect of inductive stimuli that influence M and induce the substitution of code blocks. Since β is derivable from a composition of the several α codes there is an important sense in which the α and β interaction cannot give rise to a "novel" form or to essentially "higher" levels of organisation. It is true that novel forms of differentiation could arise from externally induced mutations. But the discontinuities of genuine innovation cannot arise from any internal process unless the differentiating and developing system is constrained in some fashion that would, in the absence of cooperation, merely separate the parts (or place them in a competitive relationship, for cooperation is the converse of competition). For imprinting, the relevant constraints are social. They entail the systems of convention and tradition discussed in Chapter VII and Chapter VIII. In the developing embryo they are occasionally more tangible and depend upon the geometry that restricts growth (if you have a ball of cells, for example, that grows more rapidly at the top than at the bottom, then, geometrically, this ball must invaginate). Genuine novelty occurs when restrictions of this sort bring separately differentiated parts into informational proximity and, in the case to be cited, into physical proximity.

At a rather early stage in the development of the embryo of amphibia it is possible to distinguish a primitive or presumptive mesoderm (which arises from the so called "grey crescent" of the egg). Whilst the embryo is still a hollow ball it is also possible to distinguish primitive ectoderm (which is to be the outer layer when the ball invaginates to produce a "gastrula" and which, in the adult, gives rise to structures such as the skin). These tissues develop separately until the embryo "gastrulates". At this point the presumptive mesoderm and the ectoderm come into physical contact and the presumptive mesoderm induces the local transformation of ectoderm into neural tissue. In embryology such an occurrence is called "evocation". It is perhaps the most explicit case of novelty production and it exemplifies the type of evolutionary process which, in the abstract, has the form

- (1) Variation to produce alternatives by differing experience or external mutation
- (2) Differentiation of specifically distinct structures
- (3) Separation of these structures and further differentiation.
- (4) Approximation of distinct and separately developed structures.
- (5) Cooperative interaction to form some novel organisation.

BLANK PAGE

Difference
minimising
feedback.



which, in the present idiom, becomes:-

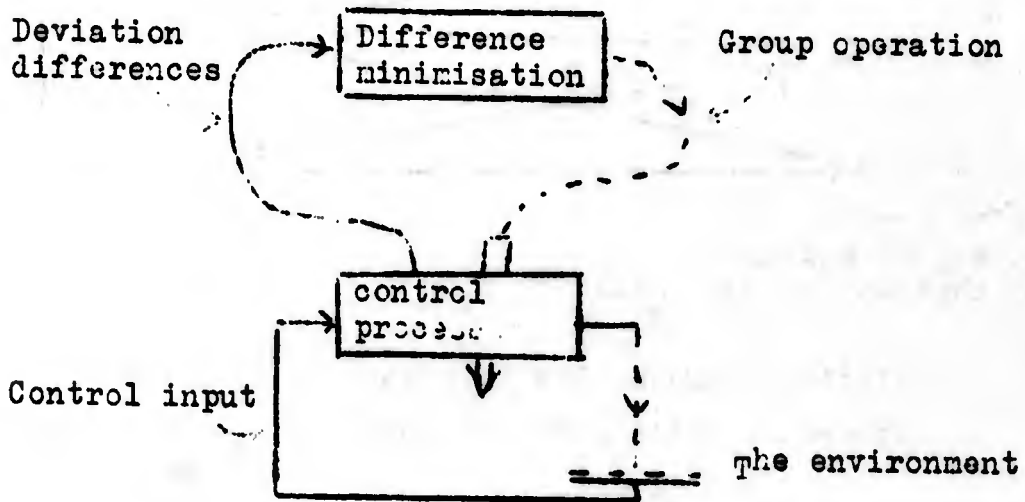




DIAGRAM 21

Reinforcement device allowing synaptic facilitation and thus change in C_i if R_i is acceptable,

Filter looking for acceptable R_i .




Simple malleable and reinforcible network with autonomous activity.



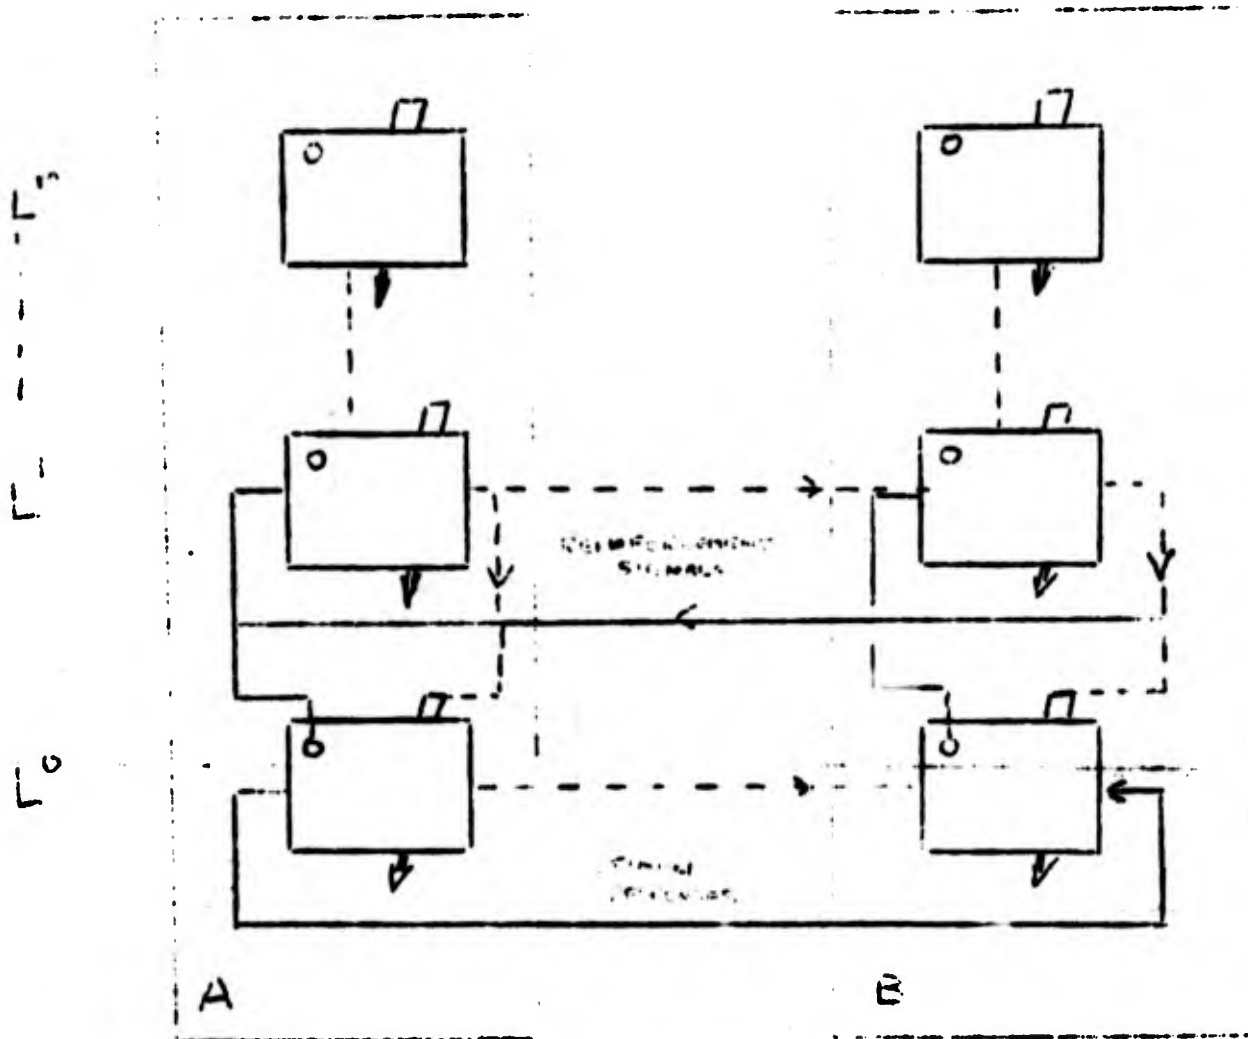
subset of malleable states C_i that induce some activity R_i .

Relation between plastic states C_i of malleable network and it's nodes of activity R_i .

DIAGRAM 25.



One of YOUNG'S proposed mechanisms for learning to avoid "pain", and to approach "pleasurable" visual stimuli.



Pair of organisms A and B conceived as hierarchically organised control systems communicating at levels of discourse

$L^0, L^1, \dots, L^m, L = 0, 1, \dots$

DIAGRAM 27.


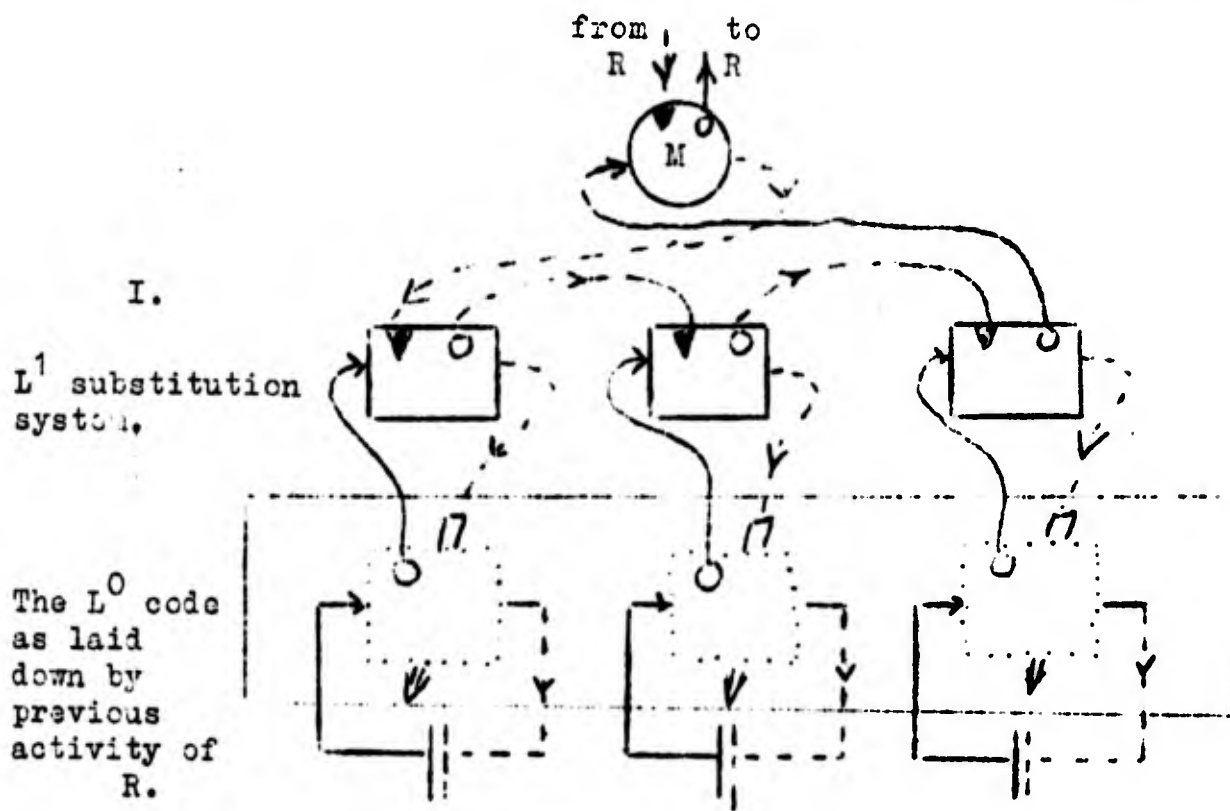


DIAGRAM 28.

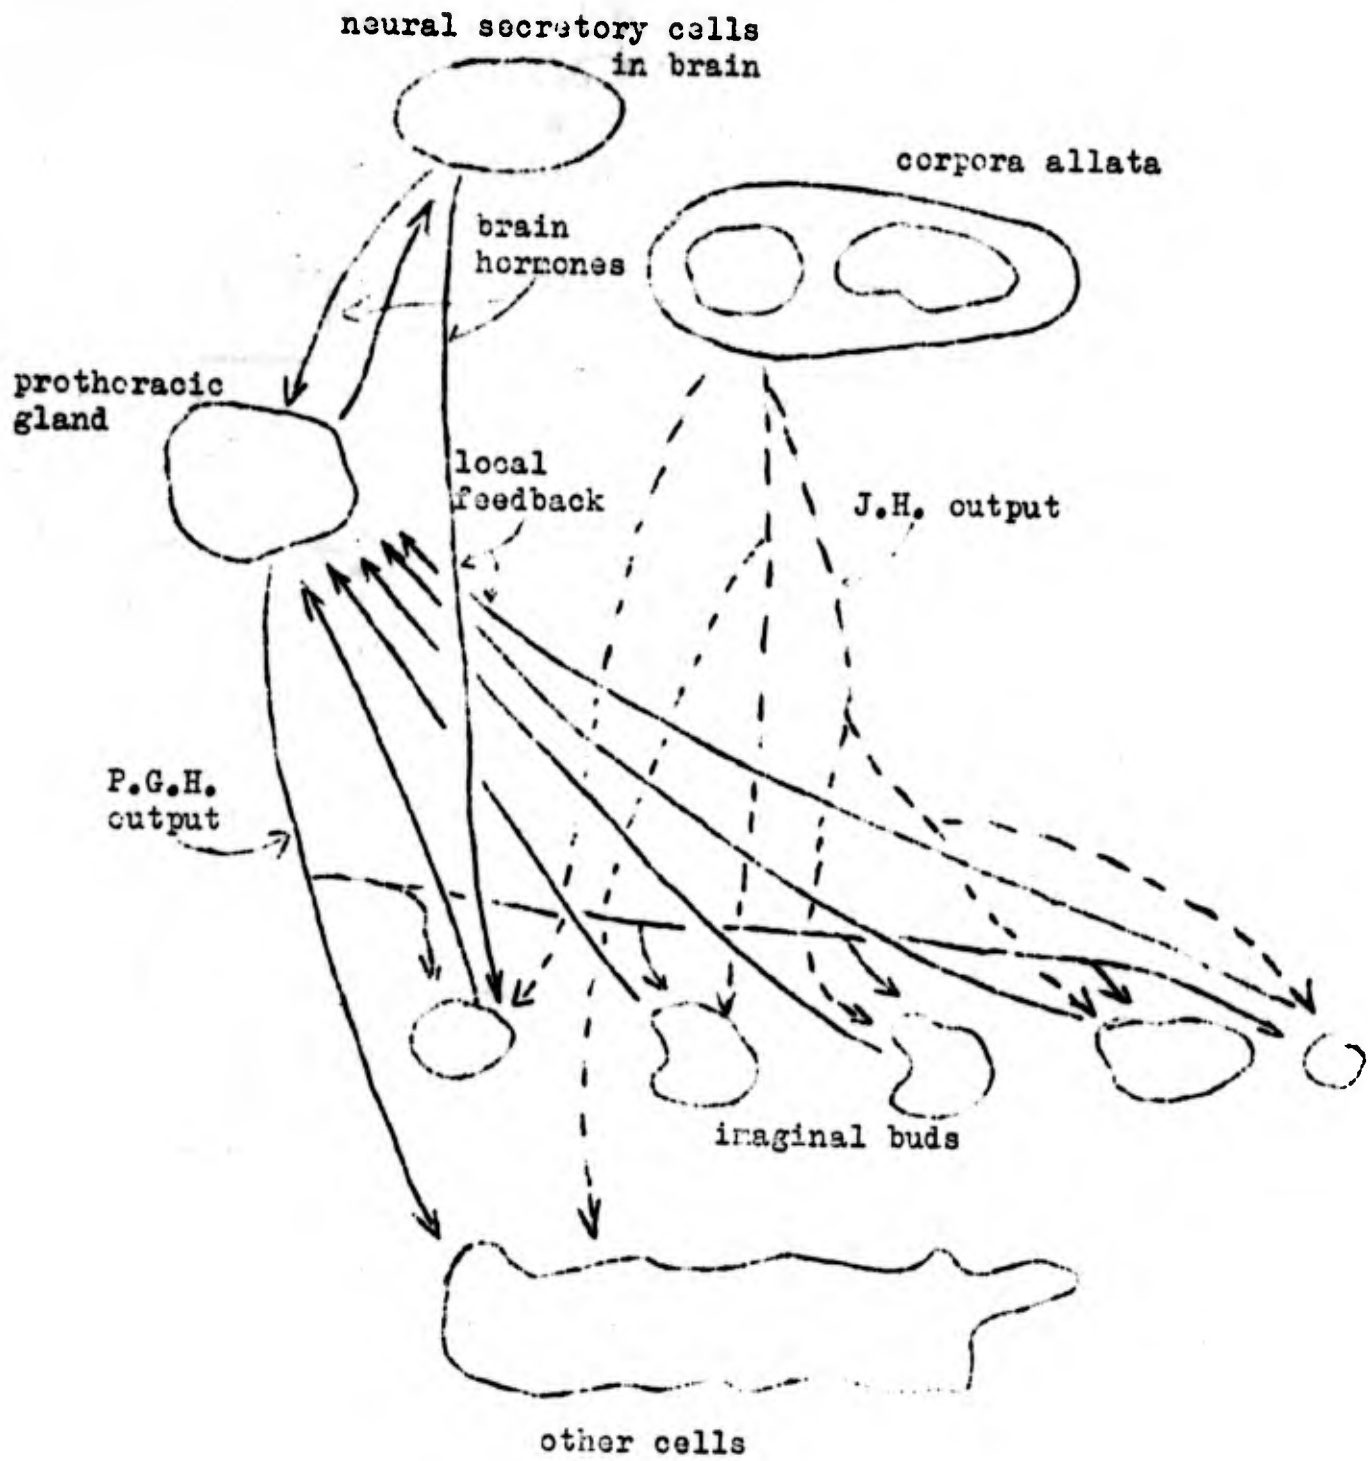


DIAGRAM 29.

The organism being imprinted.

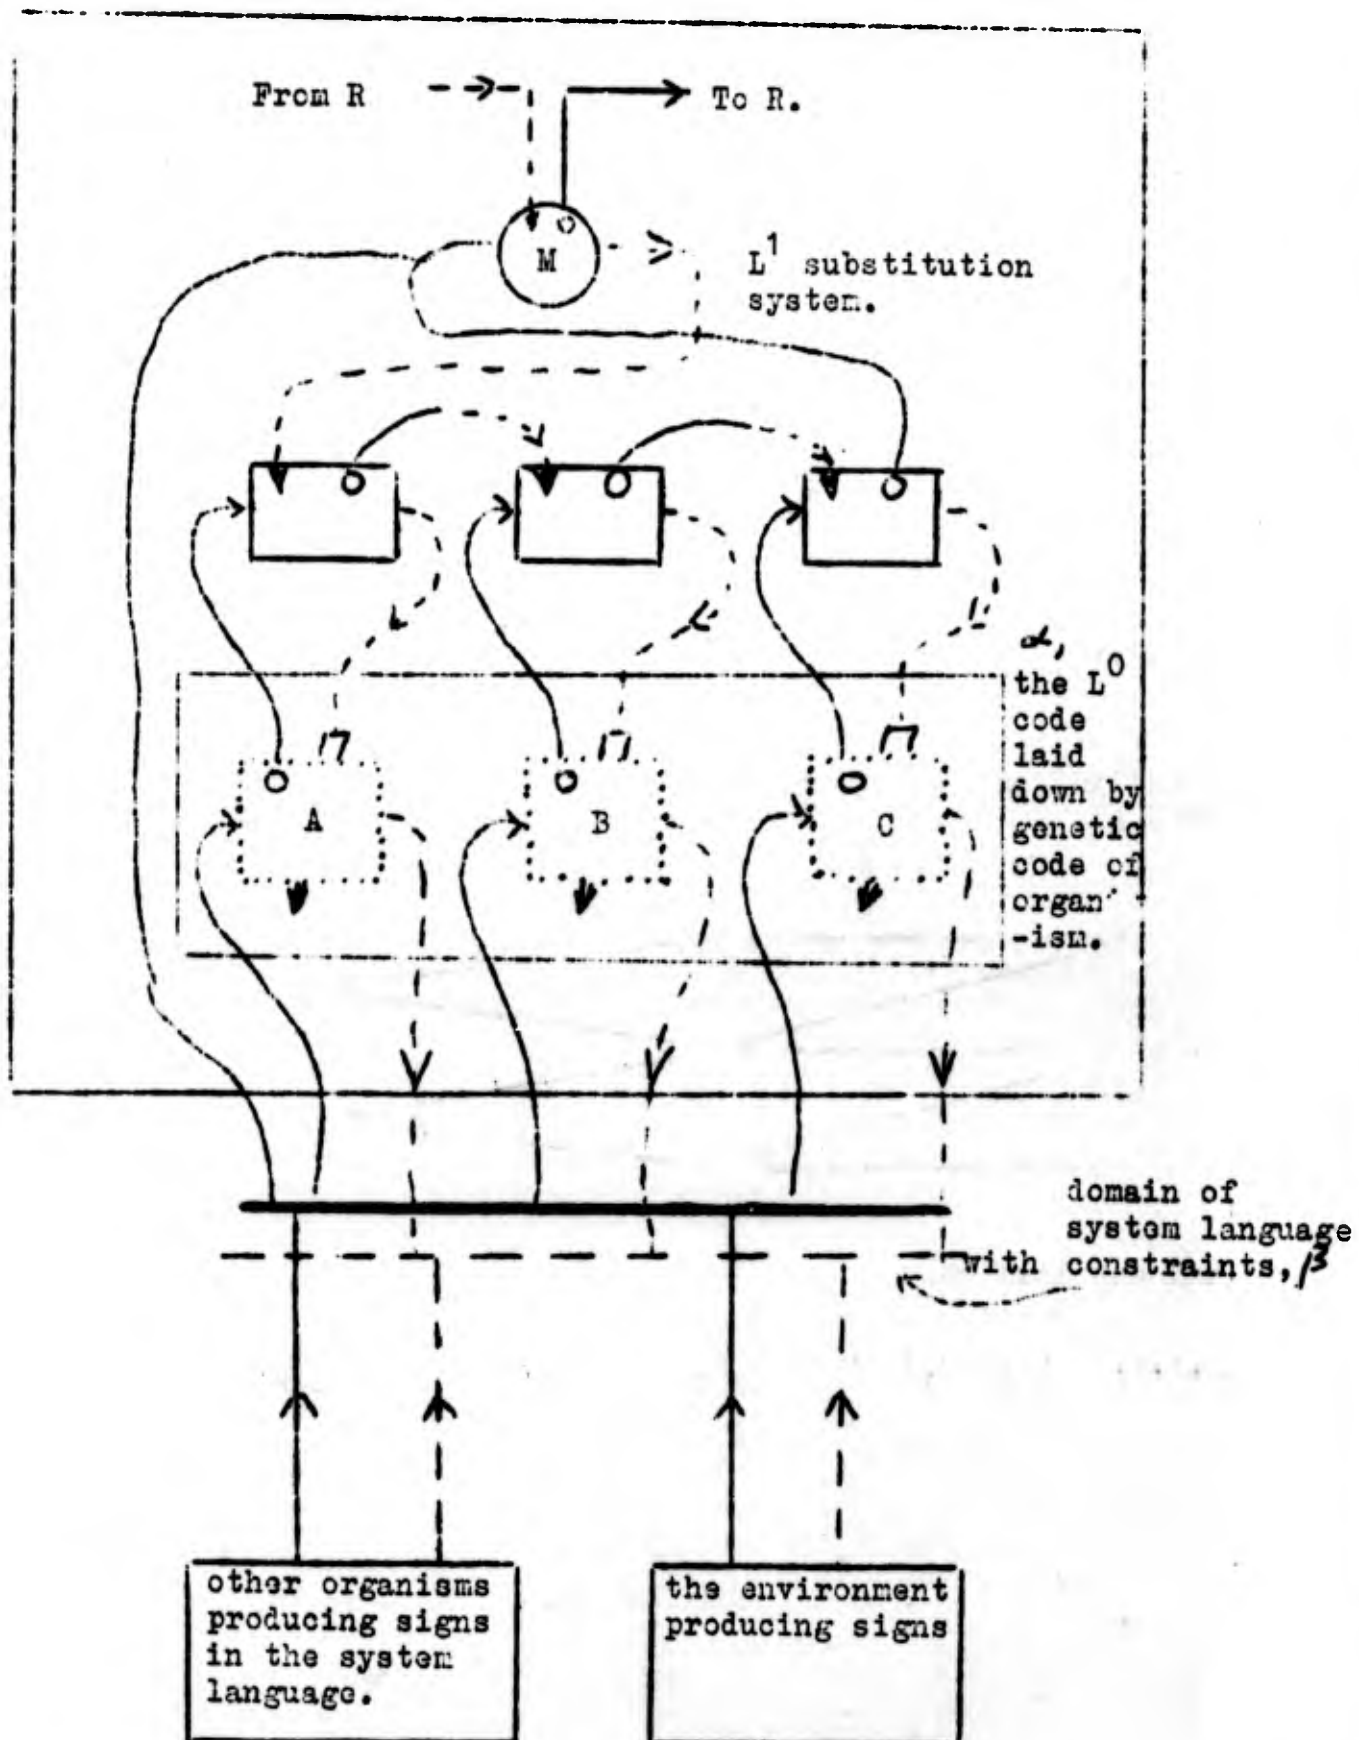


DIAGRAM 30.

α allows


(A = A₁ or A₂ or A₃ or A₄)
 (B = B₁ or B₂ or B₃ or B₄)
 (C = C₁ or C₂ only)

β allows

(A₁ or A₃ or A₄)
 (B₁ or B₂ if A₁ and
 B₂ or B₄ if A₂ and
 B₄ only if A₄.)

(C₁ if B₁
 C₁ or C₂ if B₂
 C₄ if B₄.)

Hence the substitution can proceed along certain paths, namely:-



and may lead only to :-

A, B, C, = (A₁B₁C₁) or (A₁B₂C₁) or (A₄B₂C₂) or (A₂B₂C₁) or (A₂B₂C₂)
 or (A₂B₄C₂) or (A₄B₄C₄)

DIAGRAM 31.