Summary of Research in Logical Theory of Adaptive Systems

TECHNICAL REPORT NO. AFAL-TR-65-26

March 1965

NOTICES

When Government drawings, specifications, or other data are used for any purpose other than in connection with a definitely related Government procurement operation, the United States Government thereby incurs no responsibility nor any obligation whatsoever; and the fact that the Government may have formulated, furnished, or in any way supplied the said drawings, specifications, or other data, is not to be regarded by implication or otherwise as in any manner licensing the holder or any other person or corporation, or conveying any rights or permission to manufacture, use, or sell any patented invention that may in any way be related thereto.

Qualified users may obtain copies of this report from the Defense Documentation Center.

Defense Documentation Center release to the Office of Technical Services is not authorized.

The distribution of this report is limited because of the provisions of U.S. Export Control Acts.

Copies of this report should not be returned to the Research and Technology Division unless return is required by security considerations, contractual obligations, or notice on a specific document.
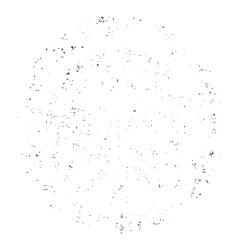
FOREWORD

This report was prepared by The University of Michigan under USAF Contract No. AF33(615)-1162. The contract was initiated under Project No. 4160, "Engineering Bionics," Task No. 416004, "Bionic Sub-System Techniques." The work was administered under the direction of the Bionics Branch, Electronic Technology Division, Air Force Avionics Laboratory, Dr. Donald E. Lewis, program monitor.

This Final Technical Report of Research on the Logical Theory of Adaptive systems covers work conducted from 1 November 1963 to 30 November 1964.

This program of research was carried out at The University of Michigan by the Logic of Computers Group, Professor Arthur W. Burks, Director. The principal investigator for this research program was John H. Holland, Associate Professor, Communication Sciences Department. Assisting Dr. Holland in this program were Marion Finley, Jr., Stephen Hedetniemi, and Carl Page, Research Assistants in Communication Sciences, and Dr. Harvey Garner, Professor of Electrical Engineering at The University of Michigan. For the highly technical details of the research effort conducted under this contract, the reader is referred to the following two Interim Engineering Reports: (1) 06114-1-T, June 1964, by John H. Holland and (2) 06114-2-T, March 1965, Holland et al.

Enqm
UMR
1440

ABSTRACT

The general objective of the program of research was to construct a logical theory of adaptive systems, through the use of automata theory and nerve net simulation. Among the sub-goals were the following: obtain a general framework for the description of learning machines in a formalized environment; investigate the formal methods by which a mechanical system can make use of regularities in its environment so as fruitfully to depart from simple enumerative behavior (i.e., seek to define efficient adaptation techniques); obtain precise mathematical characterizations of adaptations, adaptive rate, and adaptive efficiency; conduct a series of neuron net computer experiments to determine the conditions under which networks of idealized neurons would exhibit goal-directed behavior.

Specific results obtained and conclusions reached as a result of this research include the following:

a. Once the appropriate formal framework for study of adaptive systems has been determined and defined (06114-1-T), the central problem becomes the definition of construction in universal spaces. Initial steps toward this definition are described.

b. The application of a formal definition of probabilistic sequential machines to certain behavior of biological cells (the basic components of all natural adaptive systems) yields a theoretical explanation of the relation between the information coded in DNA in the chromosomes to the production of proteins, particularly enzymes, at the ribosomes of the cell.

c. From the first of a three-stage series of experiments, it was determined that under certain conditions, simple cycle-less neuron nets can correlate their output behavior with patterned inputs.

TABLE OF CONTENTS

I. A FRAMEWORK FOR THE FORMAL DESCRIPTION OF MACHINE ADAPTATION

A system of "universal embedding spaces" for automata was defined. This phase of research was necessary since a universal embedding space for automata provides a uniform common framework for comparisons of the computational and constructional complexity and efficiency of alternative particular formulations of adapting systems (comparisons which would otherwise be almost impossible to make owing to differences in presentational conventions).

To be satisfactory for the intended purpose, any growing automata system must possess the property that both the adapting mechanism and its environment must be expressible in the system; this property ensures that mathematically rigorous statements can appropriately be made about the interaction of mechanism and environment. For example, various particular systems for the expressing of automaton growth, self-duplication, etc., have been formulated (Burks,[1] Holland,[4] von Neumann,[15] Church,[3] Myhill,[14] Moore[13]). Any of these formulations can be employed to express the adapting machine and its interaction with its environment; thus each of these particular systems provides a candidate for a framework for theoretical work in adaptive systems theory. Concepts on adaptation developed in one of these systems, however, are often difficult to translate into a different system. In order to make comparisons between concepts formulated in different systems, and in particular, to compare adapting mechanisms according to their adaptive efficiency, a unified common ground of comparison, apart from particular conventions, must be found.

"Universal embedding spaces" provides a means of rigorously describing what one might mean by growth, self-description, adaptation, rate of adaptation, etc. (matters which are part of the ultimate goal of this line of research) in any of a wide range of growing automaton formulations.

A closer look at one of these particular formulations of growing automata, that of John von Neumann, will make apparent some of the properties which must be accommodated. Von Neumann had as his ultimate goal the construction of automata by automata; he therefore set up his space in such a way that complex automata could be formed by connected complexes of component automata. In addition he made his space homogeneous; at every point in the space there was a copy of the same 29-state component automaton. This meant that if his construction procedures could be carried out in one location in his space, they could be carried out in any other location.

The component 29-state automata had to possess the property that they could be combined into a complex automaton having universal computing prop-

---

erties, (i.e., it had to be possible to form a Turing machine from complexes of the underlying 29-state automata). This assured that <u>anything</u> that could be done by machine was representable, thus that no (possibly important) machine procedure had been eliminated. It is properties such as these—connected complexes, homogeneity, computational universality—which must be provided for in a universal embedding space.

Thus, according with one's theoretical or practical aims regarding various systems of growing automata for expressing the structure and behavior of adapting systems, the "universal spaces" provide a common descriptive framework.

The Interim Technical Engineering Report "Universal Embedding Spaces for Automata," (06114-1-T, June 1964, by John Holland) presents in detail the research toward this goal. A description of the particular results of this report will be given now; the detailed presentation of the theorems and their proofs will be found in the subject report.

The complexes of automata which can in general be built up from an arbitrary finite or countably infinite set of automata were characterized. In this composing of larger automata (the potential adapting automata and their environment) from component automata, certain connection properties are required. These properties are obtained by viewing automata as devices for transforming (input) sequences into (output) sequences (the transition and output functions being combined and being viewed as a functional mapping sequences into sequences). (This sequence-to-sequence viewpoint follows the work, for example, of Burks and Wright in their "Theory of Logical Nets"[2] rather than the finite string acceptance concept stemming from S. C. Kleene in his "Representation of Events in Nerve Nets and Finite Automata."[9])

By use of a composition function, either component or already "composed" automata can be combined to form yet larger automata. This composition function is so defined that certain desirable properties are retained, viz. unique solutions, finite automaton computation-universality. Computation-universality is so defined that for every finite automaton computation there is some composition of automata which can perform the computation, and which is embeddable in the universal space. The desired computation-universal property for an automaton composition requires that the composition be a countably infinite, iterated structure on a single component and with commutativity of output wire labels.

The set of homogeneous computation-universal compositions so obtained can be divided into two sub-classes according as the inputs of an automaton defined in this space be able to effect its outputs only after a delay or not.

The sub-class requiring the delay includes such important formulations of growing automata as that of von Neumann. The alternative sub-class (which includes the "iterative circuit computers" of J. Holland) possesses theoret-

ically useful, but practically rather stringent properties; it permits the (sometimes impractical) concept of an output response to a present input even though arbitrarily large numbers of connected elements may intervene. (These two sub-classes are related to the Moore,[12] and to the Mealy[11] formulations respectively, of finite automata theory.)

Again, for practical reasons, one might wish to focus on the more "realistic" universal spaces only. These would be spaces of the delay-type only, in which the further restriction is made that freedom of "fan-out" must be restricted (exponential increase of numbers of output wires is barred; see McNaughton[10]).

However, there is no way of simulating finite automata in a von Neumann array in such a way that the corresponding behaviors all occur with some constant change of time scale, $t' = kt + c$; the more complex the finite automaton simulated, the slower the simulation. On the other hand, an arbitrary finite automaton can be simulated in an iterative circuit computer, preserving not only behavioral timing (k = 1) but also details of local structural and behavioral relation.. (E.g., the simulation can reflect differences corresponding to realization in terms of the stroke function, {|}.) In fact, it can be shown such simulation is possible in iterated cellular arrays with locally finite information transfer characteristics only if there is provision for the "making" and "breaking" of non-delay paths. For development of these points the reader is referred to Ref. 5.

II. EFFICIENT ADAPTATION IN AUTOMATON SYSTEMS

Completion of the work on universal spaces provided a formal framework in which arbitrary adaptive organizations could be defined and discussed in a uniform way.[6] This enabled a start to be made on the problem of comparing various classes of adaptive systems, especially with regard to efficiency. This initial approach turned on the following ideas:

1. A common class of environments was to be specified—the tree-search environments. Each element of the class is a tree with payoff assigned to its termination (cf. von Neumann's normal form for games.[16]).

2. The output of the adaptive system at any given time was to be the "next move" specification; its input was to be the state description ("board configuration") corresponding to its current position in the tree.

3. Various organizations for adaptive systems were to be described in terms of the corresponding automata organization (growing logical nets) and then described in the universal embedding space (cf. the discussion of embedding in TDR 06114-1-T, Section 3).

4. Comparisons were to be made in the efficiency with which the various embedded systems accumulated payoff. The first notion to be used here was to compare the payoffs accumulated to time t, $\mu_T(t)$, $\mu_{T'}(t)$, by two different strategies, T and T', according to the criterion

$$\lim_{t \to \infty} \frac{\mu_T(t)}{\mu_{T'}(t)}$$

related to the notion of gambler's ruin in probability.

Several considerable difficulties were uncovered, the most important of which concerns the definition of construction in universal embedding spaces. This concept enters when the embedded system changes its organization to employ a new strategy for exploring the tree. Until this concept can be given a precise formulation for universal embedding spaces, the set of admissible strategies {T} remains undefined. This in turn leaves the domain of the above ratio criterion unspecified.

In summary, although considerable progress (and a deeper insight) has been obtained, in our objective of a formal framework for study of adaptive systems, difficult problems still remain. In our opinion the class of universal embedding spaces is the appropriate formal framework for this investiga-

tion. The first nine months of this contract were spent in giving precise
definition to this class and the resulting work is being published in a fest-
schrift for Norbert Wiener. The last three months were less productive of
concrete results, being primarily a period of discovering obstacles to further
progress and getting a clear idea of the effort required to overcome them—
central among these is the definition of construction in universal embedding
spaces.

Further research along these lines should be directed to a study of rea-
sonable definitions of construction in universal spaces followed by use of
such definitions to specify the class of admissible strategies for adaptation.
Having thus precisely specified the domain of the proposed ratio criterion,
it becomes possible to specify near-optimal efficiencies for certain tract-
able problems of adaptation and, in more difficult cases, provide necessary
conditions (guidelines) for efficient systems.

III. PROBABILISTIC SEQUENTIAL MACHINES AS MODELS
OF GROWTH AND ADAPTATION IN CELLULAR SYSTEMS

The purpose of this research was to develop a portion of the theory of
automata that would be suitable for explaining certain behavior observed in
biological cells, the basic components of all natural adaptive systems. We
wished to establish a formalism which relates the information coded in DNA in
the chromosomes of a cell to the production of various proteins, particularly
enzymes, at the ribosomes of the cell.

A relatively simple model of this process which falls within the frame-
work of probabilistic sequential machines is constructed. It differs from
many mathematical formalisms of cell operation in that it is a steady state
model, rather than a continuous diffusion model. Many details in the opera-
tion of the cell which are not yet known are lumped into individual param-
eters in the model. Some hypotheses about the operation of the idealized cell
allow it to be studied without knowing all the values of the parameters. For
example, results about the behavior of the idealized cell can be obtained
if one assumes such things as: (1) diffusion between ribosomes is symmetric,
and (2) when two or more chromosomes are active in controlling the production
of the protein, the process is independent of the order in which the informa-
tion is received, i.e., the control sequence received from the various chro-
mosomes is commutative. Assumption (2) attempts to make the mechanisms of
control independent of the random spatial arrangement of the chromosomes. Be-
fore a formal definition of probabilistic sequential machines is stated,
some examples of them will be presented.

A. EXAMPLES OF PROBABILISTIC SEQUENTIAL MACHINES

We consider here two models, one of which can be considered probabilis-
tic and one of which can be considered deterministic, although both fall
within the framework of probabilistic sequential machines. After the examples,
which are intended to clarify basic concepts, will be presented a formal defini-
tion.

Example 1. Probabilistic Internal Operation: A Slot-Machine

A simple model of a probabilistic sequential machine is a slot-machine.
The static position of the dials represents the present state of the machine.
Usually there are 20 different positions on the dial and 3 dials for a total
of 8,000 states. The input consists of putting in a coin and pulling a lever,
causing the machine to travel transiently through many states until it settles
down in one state. An output is associated with each state. Nothing (which

is associated with 0) comes out unless the dials all display the same object. In that case, some change tumbles out (which is associated with the corresponding real number) usually dependent only on the kind of object being displayed, i.e., the state of the machine. Such a machine whose output is controlled by its states is known as a "Moore machine."[12] Each state can be associated with a number between 1 and 8,000, and the output for each state can be tabulated in a column vector or 8,000 x 1 matrix. In the formalism, this column vector will be called the "output vector" and designated by the symbol "F." The output for state i will be written as "F_i."

The enormous number of distinct ways the lever can be pulled are prevented from significantly influencing the outcome by spring loading. Hence for all practical purposes there is only one kind of transition law associated with pulling the lever. If the randomness of transition of the dials caused by variable factors like dust friction, humidity, heating and small vibrations do not change over long periods of time, the probability of a transition from any state of the dials to any other can be determined experimentally to any required precision. This situation is summarized formally in the assumption for probabilistic sequential machines that the transition probabilities are stationary. Symbolizing the usual lever play of the machine by L, the transition probabilities can be tabulated in a matrix A(L), with the entry in the i'th row and j'th column (written $A(L)_{ij}$ being the probability of a transition from state i to state j via input L.

If there were no other permissible way to effect the rotation of the dials than by a pull of the lever, then the behavior of a slot-machine A could be described as a finite state Markov chain with rewards and transition matrix A(L). However, sudden small external shocks during the rotation of the dials can influence the state transitions of the machine. In order to model completely how such machines are played, we can consider a finite repeatable set of such non-standard inputs to the machine. For instance, one such input might be described as the application of a kick with a prescribed kinetic energy on a certain spot on the machine occurring 1/5 of a second after the lever is released. Symbolizing this manner of playing the machine by K, the transition matrix A(K) could be determined experimentally since the input is repeatable. A finite set of such repeatable inputs could be defined and their effects on the behavior of the machine ascertained.

To find out how strings of S and K inputs to the machine affect its operation, it is sufficient to multiply the matrices A(s) and A(K) together in the order specified by the string, e.g., if a string X is SKKSK then the transition matrix A(X) is the product $A(S) \cdot A(K) \cdot A(K) \cdot A(S) \cdot A(K)$.

Consider how the dials of the machine might be found initially. If the dials can be completely observed, the initial state of the machine is observable. In this case, in the formalism the initial state i is represented by a vector I (or a 1 x 8,000 matrix) with a 1 in the i'th component and

zeroes elsewhere. On the other hand, the dials may not be completely visible, and we may wish to specify the average behavior of a large number of machines run simultaneously or we may wish to consider the average return from playing one machine only when it is left by other players in one of a set of preferred states. In any one of these cases, I can be a stochastic vector $(I_1, \ldots, I_{8,000})$ where I_i is the probability of being in state i at time t_0.

In the general case, the next state probabilities starting with an initial state vector I and an input string X are given by $I \cdot A(X)$. Hence the expected value of output for a machine A starting with initial state distribution I and output vector F after a string X of inputs has occurred is just

$$E_A(X) \quad = \quad I \cdot A(X) \cdot F$$

which is a bilinear form in I and F with form matrix $A(X)$. The variance in output and other higher moments can be defined analogously.

Example 2. Deterministic Internal Structure: Chemical Production Cell

Suppose a chemical tank A is divided into several isolated compartments $A_1, \ldots, A_n$ by partitions which are interconnected by an electronically controlled system of pumps and valves. Suppose that there is a finite set of controls $\Sigma = 0, 1, \ldots, K-1$ and that for each control c a fixed fraction of the chemical in compartment A_i, v_{ij}^c, is pumped into compartment A_j. For all controls c in Σ, the full influence on redistribution of liquid in the tank can be described in a n x n matrix $A(c)$ with v_{ij}^c being $A(c)_{ij}$. Furthermore, suppose that the liquid being pumped between compartments is a catalyst which causes production of a desired end product in each compartment with a different efficiency, i.e., if the mass fraction of catalyst in A_i is P_i and F_i is the efficiency of A_i, then the output of end product is $P_i F_i$. Note that it is assumed that the output of the compartment depends linearly on the catalyst present.

The initial state I is an n component vector with the ith component I_i being the mass fraction of catalyst in compartment i. Note that

$$\sum_{i=1}^{n} I_i \quad = \quad 1$$

since the tank is a closed system as far as the catalyst is concerned. The distribution of mass fractions of catalyst over the compartments after a sequence of controls $X = i_1 \ldots i_m$ is just

$$I \cdot A(i_1) \cdot \ldots \cdot A(i_m) \quad = \quad I \cdot A(x) \quad .$$

8

That is, $(I \cdot A(X))_i$ is the mass fraction of catalyst in compartment i after starting with initial distribution I of catalyst fractions over compartments and the string of control inputs $X = i_1 \ldots i_m$.

The total end product from the tank is the sum of the outputs from each compartment:

$$\sum_{i=1}^{n} (I \cdot A(X))_i F_i$$

which can be written $I \cdot A(X) \cdot F$ in matrix notation. This expression has the same form as the expectation of output for the probabilistic slot-machine, but there are no overt probabilities involved here. The mass fractions of catalyst play the same role as the probabilities in the first example. However, the output will still be written like an expectation as $E_A(X)$.

The total end product accumulated, T_X, for the string of controls X from time t_0 to time $t_0 + m$ is given by adding the output from each substring, i.e.,

$$T_X = E_A(i_1) + E_A(i_1 i_2) + \ldots + E(i_1 i_2 \ldots i_m)$$

B. FORMAL DEFINITION OF PROBABILISTIC SEQUENTIAL MACHINE

By a probabilistic sequential machine is meant a system which satisfies one of the following two definitions:

Definition 1.1: A (Moore-type) probabilistic sequential machine.
A is a system $A = < n, I, S, \Sigma, A(0), \ldots, A(k-1), F, O >$
where

      n: a natural number, the number of states

      I: an n-dimensional stochastic vector, the initial state vector

      S: set of state vectors = $\{S_1 = (1,0,\ldots,0), \ldots, S_n = (0,\ldots,0,1)\}$

      Σ: alphabet set $\Sigma = \{0,1,2,\ldots,k-1\}$

    A(i): i = 0, 1, ..., k-1 n x n switching matrix for input symbol i. $A(i)_{\ell m}$ is the probability of a transition from state ℓ to state m via symbol i.

      F: output vector, an n-dimensional column vector whose entries are real numbers

      O: output function $O(S_i) = S_i \times F = F_i : S_i \in S$

Definition 1.2: A (Mealy-type) probabilistic sequential machine.
$A = < n, I, S, \Sigma, A(0), \ldots, A(k-1), W, P >$

where n, I, S, Σ, A(0),...,A(k-1) are as in 1.1 and where the output function P satisfies

$$P(S_i,j) \quad = \quad W_{ij}S_i \, \epsilon \, S, \quad j \, \epsilon \, \Sigma$$

It is an easy matter to show that Definition 1.1 and 1.2 are equivalent in the following sense: For every Moore-type probabilistic sequential machine there is a Mealy-type sequential machine whose output is the same random variable over each input and vice versa. Consequently, we will be concerned only with the properties of Moore-type probabilistic sequential machines, which from now on will be called probabilistic sequential machines.

C. APPLICATIONS TO BIOLOGICAL CELLS

Comparing examples 1 and 2 to the formal definition we see that both indeed are probabilistic sequential machines. We now consider a tentative model of the production of protein by the ribosomes of a biological cell which is based upon example 2 and hence views the cell as a probabilistic sequential machine. We regard the cell as the tank and the compartments as partitioned areas in the cytoplasm around ribosomes. Instead of electronic pumps and valves, chemicals cause changes in the isolation of the compartments, perhaps by changing the permeabilities of membranes in the cytoplasm. The input control alphabet, controlling the chemicals affecting the isolation of the ribosomes, are distinct sequences on the chromosomes which cause steady state changes in the distribution of the catalyst. A transition matrix might be experimentally determined for each sequence. We suppose in addition that the catalyst linearly affects the production of some protein at the ribosomes and that the ribosomes have different efficiencies. Hence the output of the cell of the protein, which may be an enzyme, can be changed from one control time to the next by redistribution of the catalyst among the ribosomes. Whether these suppositions are correct for some protein will have to be determined experimentally. However, from them can be deduced certain aspects of the behavior of the cell which may be more easily observed.

At present the theoretical framework of probabilistic sequential machines (P.S.M.) is being used to study two important phenomena in the behavior of cells. The first phenomena is that of mosaics in the chromosomes of cells, i.e., the influence on the protein production of the cell by one, two, or three extra chromosomes due to a hereditary accident. P.S.M. theory has been used to derive inequalities between the amount of protein produced by normal cells and mosaic cells, subject to the assumptions (1) and (2) of the introduction.

The second phenomena concerns the underlying causes of cancerous behavior of cells. Presumably such behavior is caused by alterations of the production of enzymes. There is a longstanding dispute among experts in carcinogenesis as to whether the behavior of the cell is more severely affected by ruptures

in the cell wall or by changes in the chromosomes.

Ruptures of the cell walls undoubtedly influence the isolation of the ribosomes. In the P.S.M. model, ruptures are considered to cause merging of states of the machine, analogous to ribosomes being forced into the same chemical neighborhood. There is uncertainty in the model as to what efficiency to assign to the state corresponding to a rupture area containing two or more ribosomes based upon the initial efficiencies of the ribosomes. However, certain plausible candidates have been tried in lieu of any experimental evidence. The study of the behavior of the machine after certain of its states have been merged has its foundations well established mathematically in the study of deflated matrices. Research will be continued on the influence of different kinds of state merging on properties of the behavior.

If definitive results about the output could be obtained from the study of merged states, they could be compared to outputs obtained by altering the input sequences in ways reflecting how chromosomes are known to be defective. Then for a given machine or for a given class of machines one could establish (perhaps by simulation) which kind of change, altering of input symbol or merging of states affects the output more drastically. Given a valid model of the protein production of the cell, this would settle the dispute as to whether ruptures in the cell walls or defective chromosomes are the more important factors in carcinogenesis.

The two major problems in biology at which this research is aimed cannot be resolved until both a large amount of theoretical effort and relevant experimental data are obtained. It was the intent of this research to advance the theory as far as possible in light of what is known while suggesting measurements to be made and plausible hypotheses to be checked.

D. SUGGESTIONS FOR FUTURE RESEARCH

This research has suggested two auxiliary areas of investigation. As we have pointed out, the expected value of output for a string X for a probabilistic sequential machine is a bilinear form in I and F with form matrix A(X). This nice mathematical description occurs because we assume the output from a state is linearly dependent on the probability of being in that state. A more general approach would be to assume the output from a state is some polynomial function of the probability of being in that state. However, the expected value is no longer a bilinear form and standard matrix methods have not been developed to study such formulations. A theoretical study of such structures might be of value to other fields in which bilinear forms are used for linear approximation because of the lack of something better.

A more radical change in the model for a cell suggests another area of investigation in matrix theory. Suppose each "symbol matrix" that is applied because of the activity of the chromosome controls the diffusion for several

time steps, so that the proper model for the diffusion process involves powers of the "symbol matrices." In fact we may assume that a symbol acts on the distribution of concentrations until it drives the vector of concentrations into an eigenvector of the symbol matrix which is a stable distribution of concentrations, i.e., the diffusion process has reached a steady state.

The mathematical question is the following: Given the K symbol matrices $A(0),\ldots,A(K-1)$ one can compute their stable distributions or eigenvectors (vectors such that $V\,A(i) = c_u V$ for some constant c_u) and usually there will be n such eigenvectors for each n x n matrix. Some will be stochastic and hence possible distributions of catalysts. But given an arbitrary stochastic eigenvector of the machine, which other stochastic eigenvectors can be reached from it by application of the symbol matrices in the manner described above? That is, what sequence of steady state operation is possible in such a system with the answer being given in some computational manner from the symbol matrices.

IV. NEURAL NETWORK EXPERIMENTS

The object of this study was to devise a series of experiments to test the self-organizing capability of large neural network systems in realistic learning situations. The theoretical basis for the work is contained in the Hebb-Milner cell-assembly theory, especially as developed later by J. H. Holland and J. W. Crichton.[8] The series of experiments was divided into three stages: the first stage deals entirely with simple, cycleless nets and is designed to allow determination of the basic network parameters associated with the threshold, fatigue, and synapse-value functions. The central idea of this sub-series of experiments was to determine whether a neuron of the model could distinguish patterned inputs from unpatterned ones by correlating its output behavior with the patterned inputs. This property must be present to ensure the self-organizing capability ultimately desired, i.e., the formation and growth of cell-assemblies. The second stage would continue with the basic network configuration of the first, but introduce gradually more and more feedback among the neurons, to the point where some of the theoretical observations made by Crichton concerning the behavior of cell assemblies could be tested, as well as such things as the alternation of activity of cell-assemblies required by Hebb's basic theory. The third stage would deal with learning experiments as such, testing the adequacy of the cell-assembly hypothesis in learning. A sub-series of experiments has been planned for this third stage.

The first stage only has been completed thus far and the correlation hypothesis was shown to hold under certain conditions. A number of experiments were performed using simple, cycle-less nets in an attempt to derive the correct functional forms of the threshold, fatigue, and synapse-value functions, all of these being essential parameters of the models. On the basis of the experimental results, a derivation was obtained for the threshold curve and the general forms of the fatigue and synapse-value functions. These last two functions, it turns out, must be hysteresis functions of a certain type whereas the former, the threshold function, must be a hyperbolic function of the recovery state.

The basic theory of Crichton (see J. W. Crichton, Doctoral Dissertation, University of Michigan, 1964) concerning trends in synapse-level changes was demonstrated by the experiments of the first stage.

For the future, it remains to test the functional forms of the network functions mentioned above, then to begin the second stage of experimentation, where cycles are introduced into the networks. It is hoped that this second series of experiments can be concluded with a test of Hebb's hypothesis of alternation of cell-assembly activity.

13

V. FURTHER RESEARCH

The following research goals and directions are suggested as possibly fruitful areas of further investigation.

A. EFFICIENT ADAPTATION IN AUTOMATON SYSTEMS

1. Conduct a study of reasonable definitions of construction in universal spaces.

2. Specify the class of admissible strategies for adaptation, using the definitions studied in 1.

3. Attempt to specify (through the application of 1 and 2) the near-optimal strategies for tractable problems of adaptation. In more difficult cases, 1 and 2 can be utilized to provide necessary conditions (guidelines) for efficient systems.

B. PROBABILISTIC SEQUENTIAL MACHINES AS MODELS OF GROWTH AND ADAPTATION IN CELLULAR SYSTEMS

1. Investigate the mathematical structure of probabilistic sequential machines whose output from a state is a polynomial function of the probability of being in that state.

2. Characterize the sequence of eigenvectors which can be obtained by arbitrarily numerous applications of the symbol matrices in order to construct a steady state diffusion model of the cell.

C. NEURAL NETWORK EXPERIMENTS

1. Test the functional forms of the network functions discussed in Section IV of this report.

2. Conduct a series of experiments introducing cycles into the networks.

REFERENCES

1. A. W. Burks (1960): Computation, behavior and structure in fixed and growing automata. Self-Organizing Systems, Marshall Yovits and Scott Cameron, Editors, New York, Pergamon Press (pp. 282-311).

2. A. W. Burks, and J. B. Wright (1953): Theory of logical nets. Proceedings of the Institute for Radio Engineers, 41, 1357-1365.

3. A. Church (1957): Application of recursive arithmetic to the problem of circuit synthesis. Summaries, Summer Institute for Symbolic Logic, Cornell University, Institute for Defense Analysis (pp. 3-50).

4. J. H. Holland (1960): Cycles in logical nets. Journal of the Franklin Institute, 270, 3, 202-226.

5. J. H. Holland (1960): Iterative circuit computers. Proc. 1960 Western Joint Computer Conference, pp. 259-265.

6. J. H. Holland (1964): Universal Embedding Spaces for Automata. To appear in a festschrift honoring Norbert Wiener, and distributed as Technical Report 06114-1-T, June 1964.

7. J. H. Holland (1964): Universal Spaces: A Basis for Studies of Adaptation. To appear in the Proceedings of the International School of Physics, Course on Automata Theory; Academic Press.

8. J. H. Holland and W. Crichton (1959): A New Method of Simulating the Central Nervous System Using an Automatic Digital Computer. Report No. 2144-1195-M, Willow Run Laboratories, The University of Michigan.

9. S. C. Kleene (1956): Representation of events in nerve nets and finite automata. Automata Studies, C. E. Shannon and J. McCarthy, Editors. Princeton, New Jersey, Princeton University Press (pp. 3-41).

10. R. McNaughton (1962): On nets made up of badly timed elements, I. Summer Conference Notes on Parallel Computers, Automata and Adaptive Systems, Ann Arbor, The University of Michigan.

11. G. H. Mealy (1955): A method for synthesizing sequential circuits. The Bell System Technical Journal, Vol. 34, pp. 1045-1079.

12. E. F. Moore (1956): Gedanken-experiments on sequential machines. Automata Studies, C. E. Shannon and J. McCarthy, Editors, Princeton, New Jersey, Princeton University Press (pp. 129-154).

REFERENCES (Concluded)

13. E. F. Moore (1962): Machine models of self-reproduction. Mathematical problems in the biological sciences, Providence, American Mathematical Society, pp. 17-33.

14. J. Myhill (1962): Self-reproducing automata. Summer Conference Notes on Parallel Computers, Automata and Adaptive Systems, Ann Arbor, The University of Michigan.

15. J. von Neumann (unpublished): The theory of automata: construction, reproduction, homogeneity.

16. J. von Neumann (1947): Theory of Games and Economic Behavior, Princeton, New Jersey, Princeton University Press.