A REVIEW OF
THE FIELD OF ARTIFICIAL INTELLIGENCE
AND
ITS POSSIBLE APPLICATIONS TO NASA OBJECTIVES

Author: M. J. Pedelty

The American University, Washington, D. C.
(School of Government and Public Administration)

FINAL REPORT
Under Contract No. NASr-229
February, 1965

To
NATIONAL AERONAUTICS AND SPACE ADMINISTRATION
Washington, D. C.

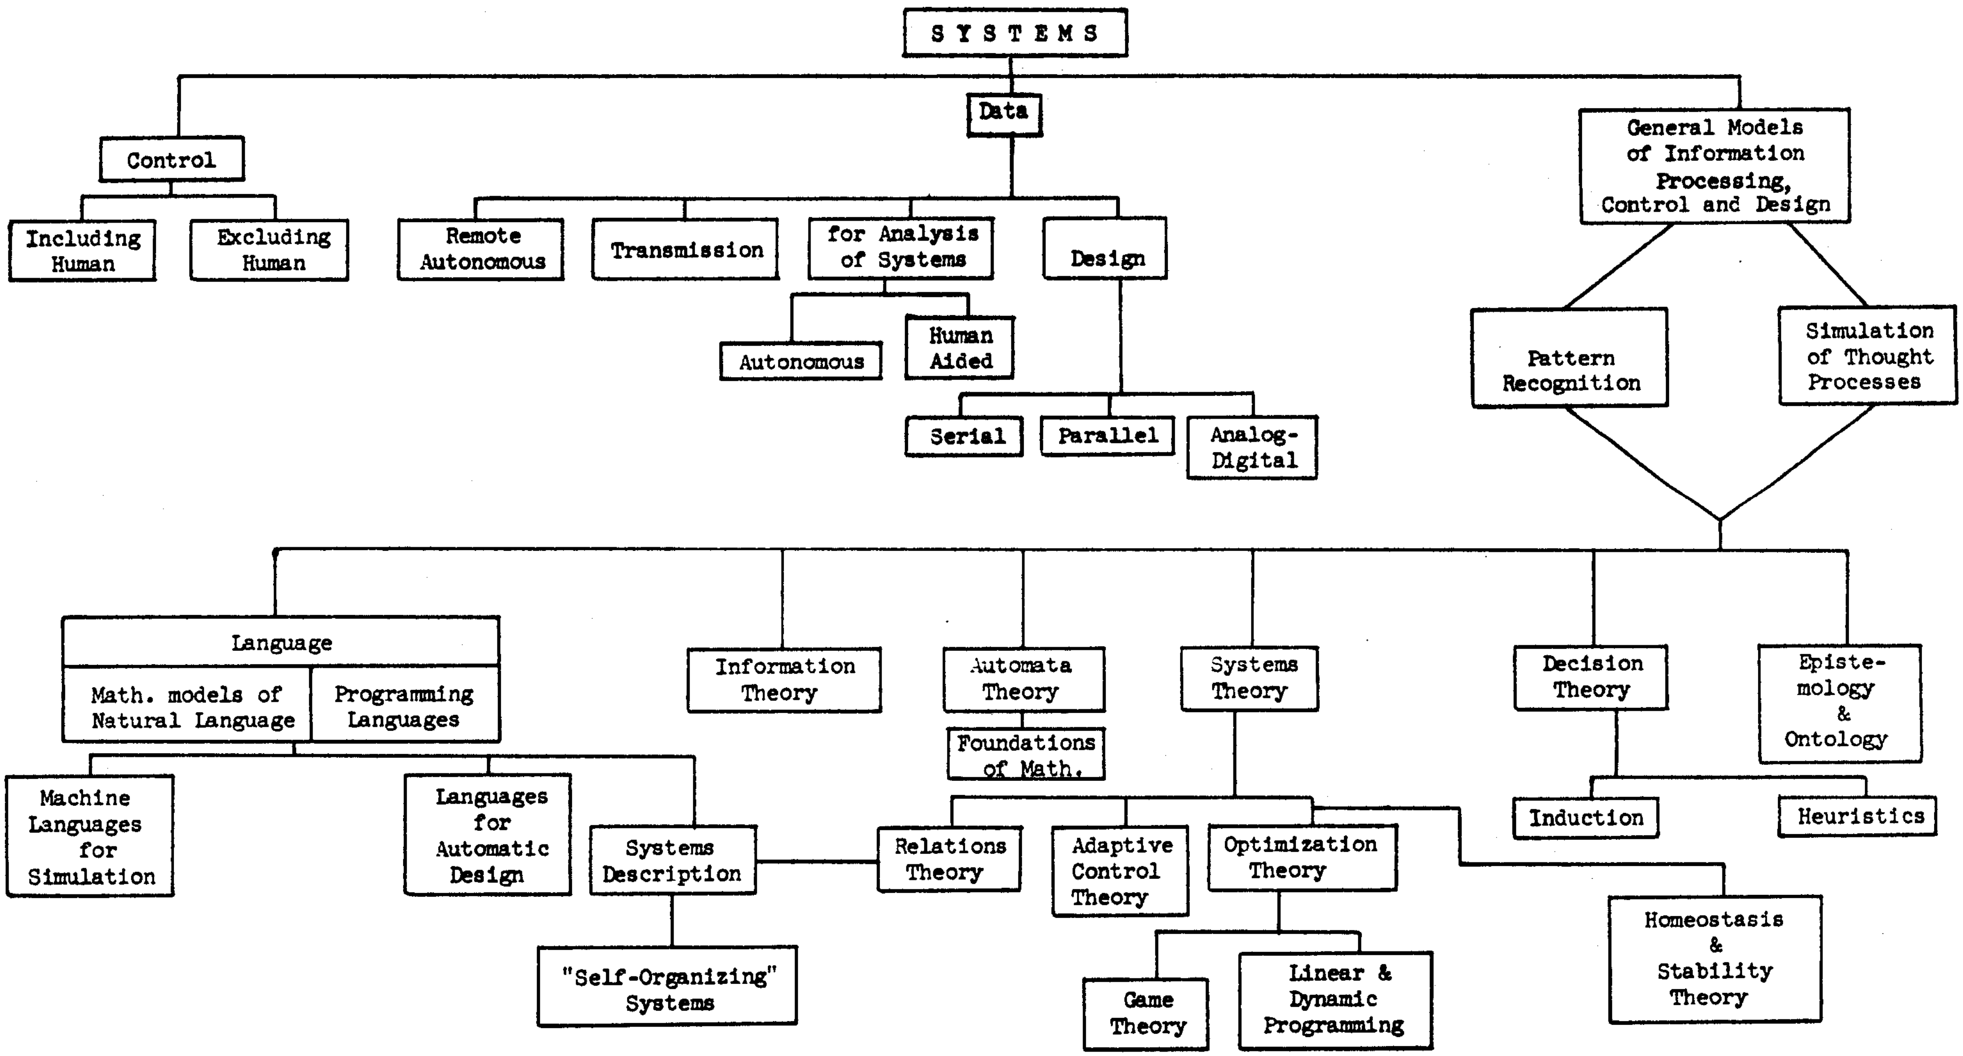

SYSTEMS
Control
Including Human
Excluding Human
Data
Remote Autonomous
Transmission
for Analysis of Systems
Autonomous
Human Aided
Design
Serial
Parallel
Analog-Digital
General Models of Information Processing, Control and Design
Pattern Recognition
Simulation of Thought Processes
Language
Math. models of Natural Language
Programming Languages
Machine Languages for Simulation
Languages for Automatic Design
Systems Description
"Self-Organizing" Systems
Information Theory
Automata Theory
Foundations of Math.
Systems Theory
Relations Theory
Adaptive Control Theory
Optimization Theory
Game Theory
Linear & Dynamic Programming
Homeostasis & Stability Theory
Decision Theory
Induction
Heuristics
Episte-mology & Ontology

TABLE OF CONTENTS

1. CONTROL SYSTEMS

1.1 What has been done

The domain of this report is the design of control systems and of data-gathering, data-analyzing systems. It also includes a discussion of formal tools and techniques which are common. In reviewing past achievements, no attempt will be made to be comprehensive. Rather, certain topics will be highlighted to the extent that they bear on problems that have still to be faced.

Control systems normally fall into one or the other of two categories: closed-loop systems without a human in the loop, and with relatively constant command inputs; and those systems containing a human in the loop. In the beginning, only those systems which were thoroughly understood, and for which simple performance criteria could be selected and implemented, were designed to operate without human feedback. In all other systems, the human was still required to supply a considerable measure of judgment, decision-making and adaptation. The history of automatic control has consisted of the gradual taking-over of more and more of these human functions by machinery designed to approximate one or more aspects of the human's performance.

By and large, evolution of automatic systems has involved repetition of three basic steps: (1) obtain an operating system with a human in the loop, (2) decide what the human does to control the system, (3) design machinery to duplicate the human control activity. If this latter duplication is sufficiently precise and reliable the human may then be removed from the system. If this criterion has not been met, the human must be retained, but, in this case, if the automatic portion of the system is doing a reasonable share of the work, the human is at least relieved of some of the tedium, and may better divide his attention between several tasks. An aircraft autopilot is an example of this. By providing stability in the roll, pitch and heading axes (and, possibly, the ability to follow a radio "beam"), the control system relieves the pilot of a great deal of work and enables him to concentrate on flight-planning and navigation. At present, however, the pilot is still an essential part of the total control system and must observe, predict and avoid other aircraft, severe weather and obstructions. The solution chosen is typical: the problem is divided into routine and non-routine tasks. The routine tasks are assigned to the automatic portion of the system. As better sensors and actuators become available, this portion may take over more of the pilot's duties (e.g., "beam" following).

1.2 What problems remain

One of the difficulties arising during the design phase is that of obtaining an exact description of the tasks actually being performed by the pilot. Although the prediction for learning-machine development leads us to believe that much task description duplication can be automated, the present state of the art of controllers (even adaptive controllers) requires minute specification of the variables involved and their interrelationships.

For taking over the tasks of a human controller, the ideal learning machine should be responsive to all aspects of the controller's sensory environment and it should be able to correlate these with his subsequent actions, correct action prediction being appropriately rewarded. Unfortunately, it is too naive to expect any loosely organized black-box to learn to make these predictions from scratch. More research is required into partitioning control tasks into temporally or informationally "disjoint" subsets and taking advantage of this information in designing the "black-box". [Sprague III, 1964] [Hartmanis, 1961]

It is probable that there should be more simulation of human adaptation and decision-making, especially in multivariable and non-linear situations so as to obtain a better understanding of what a human controller actually does. It seems likely that hybrid computing equipment would prove valuable for such simulation.

On a more theoretical plane, studies should be made on the effect of language (including mathematics) on our ability to describe any system. This task relates to fundamental problems in ontology, epistemology and the theory of relations. [MacKay, 1956] [Günther, 1962]

Simulation of physical systems needs to be stepped-up and special experiments devised for systems which do not appear amenable to earthbound simulation. On-board simulators may be needed, especially for tasks such as re-entry after a very long voyage. Research must be started on measuring the deterioration of human performance with lack of practice in this context.

Much work remains to be done on learning machines and learning control systems. Purdue University has done some good definitive work on the latter, including the use of adaptive threshold elements ("neurons") in a bang-bang servo system. The weights of these elements adapt to define a new switching boundary. [Luisi, 1962] The work appears to be realistic and promising. The coupling of learning machines to real environments, via available sensors, has not been adequately studied. It should be stressed as a part of future research.

Under this contract, a very careful study has been made of the status of Automata Theory, including a two-week visit to the University of Michigan. As far as can be determined, much of the material already available in this field cannot be directly applied to the problem being discussed here. Automata theorists have often tended to disregard the finite and the considerations that are imposed by systems which operate in real time. There are, of course, exceptions to this, and the work of Hartmanis and of Krohn and Rhodes springs at once to mind. (Although I am not personally familiar with it, I believe A. Rosenfeld of the University of Maryland has also done valuable work in this area, especially as it relates to pattern recognition.) Considerable advances in the application of Automata Theory can only be expected when more Automata theorists and algebraists become interested in finite systems which have to converge to solutions in real time. (This is further discussed in section 3.2 of this report.)

In 1959, I initiated a study to show how the adaptive threshold element ("plastic neuron") could form a self-organizing pattern recognition scheme. The system was trained by exposure to "square" images, which resulted in its ability to respond when "nearly square" images were present. (This result is, in some ways, analogous to Pavlovian conditioning.)

It was hypothesised that the "square" could be a perfect return from a radar target: the noisy "nearly square", an imperfect return. [Pedelty, 1963a] This particular study was a computer simulation. Similar systems have been built at Cornell (Rosenblatt, et al.), Stanford (Widrow, et al.) [Widrow, 1963] and S.R.I. (A. E. Brain). A group, which includes Diane M. Ramsey of Douglas/Astropower, is also active in this area. Both construction of prototypes and further simulation studies are required. (For a set of recent papers in this area, see ref. [B]).

2. DATA SYSTEMS

One of the more exciting research aspects of this entire area is the design of data-gathering, data-analyzing systems. This research has hitherto been divided into systems in which man provided the final decision-making function (which we shall call man-machine systems) and those in which completely autonomous activity would be the final research goal.

There exist missions (such as extended lunar surface exploration) for which it is difficult or impossible to provide life-support systems of more than limited duration. Such missions should be carried out by remote autonomous vehicles capable of automatic sample-taking. Such samples would then either be subject to automatic on-board analysis and data-transmission or they would be stored and labelled for return to base (e.g., a manned vehicle). The vehicles would be required to overcome natural obstacles and perform terrain exploration and survey, correlated with geological sampling. Various traction devices, including legs, should be investigated.

At various stages in an "autonomous vehicle -- manned base -- earth base" system, data must be compressed. The degree of this would be related to (1) quantity of data, (2) transmission delays due to physical distance (in light-minutes). Relatively routine data reporting and command signals could have preprogrammed compression applied. However, the more significant unexpected results would have to be sifted according to very general criteria defining the mission objectives.

These criteria would probably have to be built into self-organizing pattern recognition devices of a rather sophisticated nature. Such devices have only been studied in a rather superficial way and much more simulation and construction of prototypes will be required.

Analysis of data and of proposed man-built systems can also be enhanced by special devices and heuristic programming [D]. Studies such as "Sketchpad" [Sutherland, 1963], and B.B.N.'s "Cyclops" (Semi-Automatic System for Partitioning Complex Data Spaces) are examples of man-machine systems offering powerful concept formation enhancement. [Hunt, et al., 1963] A system similar to B.B.N.'s but using the author's "p-machine" components seems equally feasible. [Pedelty, 1964] Again, much development work is required here.

3. GENERAL THEORETICAL CONSIDERATIONS

This report has, so far, dealt with two practical types of interesting systems: control systems and data systems. The conception and design of such systems is coming to depend increasingly on a large body of theoretical knowledge applicable to all cybernetic systems. All such systems involve some kind of pattern recognition process (not necessarily of visual images, of course) and the most sophisticated are inextricably related to efforts to simulate the processes of the human brain. This is, indeed, the true domain of "artificial intelligence". There are some half-dozen distinct but closely related aspects of this problem.

3.1 Language theory (including mathematics)

This topic includes mathematical models of natural languages [Simmons, 1965] as well as machine language synthesis. [Dimsdale, et al., 1964] Both topics are closely related to automata theory [Friedman, 1962] [Chomsky, 1959]. An aspect which is still much neglected is the relationship between a language and the various concepts and meanings which are thereby available to its user. Although simulation language design is beginning to attack this problem, the real roots have yet to be properly understood. As a crude example, consider the difficulty of imparting to another the feeling of an emotional response or of a headache. Some Western observers feel that Eastern languages are more adept at this. Consider also the relationship between speech and conscious thought. Would it ever be possible for a person's semi-conscious mental activity to appear in recognizable form in parallel on many output channels? Is Aristotelian logic causing "mental blocks"? In designing intelligent machines, it is important to model not only the conscious processes of the human brain but also the great unconscious activity. [Bernhard, 1962] Much organization of knowledge, pattern recognition and trial response generation and testing appear to form more or less unconscious activities. The true value of these has not always been recognized in simulation studies.

Of particular interest is the problem of formal analysis and synthesis of languages for system description, (including, especially, "self-organizing" systems). Most systems can be fully described (from an information sciences point of view) in terms of: (1) their topology, (2) their activity. (Both of these can be represented by "relations".) As an example, the connections of a neural-glial system are its topology, while its activity consists of electrical and chemical events within that topology. It is the opinion of at least one physiologist [Lettvin, 1964] that presently available mathematics is inadequate to represent neural activity. An attempt has been made by the author to model learning and adaptation in terms of a change in mapping: such mapping being from an input "alphabet" to an output "alphabet" (or from a "stimulus set" to a "response set"). [Pedelty, 1964]

3.2 Automata Theory

This body of knowledge has already been mentioned several times in this report. It is my opinion that studies of adaptive and self-organizing systems must depend increasingly on formal automata theory and that this can only be the case in proportion to the extent to which the latter concerns itself with the finite and with considerations imposed by dynamic environments which demand that computation be completed within a fixed time. It is a curious fact, however, that automata theorists have concerned themselves very little with adaptive systems. [Holland, 1962] Indeed, the concept of teleology rarely enters these studies. On the other hand, certain new results in the foundations of mathematics and in language theory appear to depend heavily on automata theory. (Some interesting work has been started at the University of Michigan on the subject of "universal" automata embedded in two dimensional arrays. [Holland, 1962] These results should be valuable for deciding on primitives for systems simulation languages. [Holland, 1964])

3.3 Information theory

This has so far yielded a few useful results relating to coding, error-reducing redundancy, and channel capacity. In the area of computation two outstanding (but now outmoded) papers are: [Cowan, 1962] [Ray-Chaudhuri, 1961].

For self-organizing system theory, no very significant contribution will be made by "information theory" until it can take account of the significance of a piece of information relative to a machine's goals and action requirements. [Pedelty, 1963b, Ch. 4 and p. 13]. The design of "attention controls" is, of course, inextricably bound to our concept of "meaningful information".

3.4 Systems theory [Wymore, 1964]

This is a subject area without sharp boundaries and it includes homeostasis and stabilization theory; optimization theory (which can include linear and dynamic programming and game theory); adaptive control theory; the theory of relations (as applied mathematics); and system description.

All these sub-areas are fairly well advanced with the exception of the last two. In particular, under system description, much more sophisticated indices of performance are needed [Westcott, 1962] [Sarture, *et al.*, 1963] From an information theory point of view, the entropies of most such indices are pathetic, considering the information content of the controlled system. The selection of these indices has, in the past, often been influenced by the extent to which they were tractable from a design viewpoint. It is certainly to be hoped that, with greater man-machine interaction possible in the computation phase of design (*e.g.*, "sketchpad" and its derivatives) the "tractibility" of a performance index will no longer be an important consideration.

3.5 Epistemology and Ontology

These are subjects which have been grossly neglected in relation to systems synthesis, perhaps because of their very fundamental nature. However, at least two pertinent reports exist. [Gilstrap and Pedelty, 1963] [Wallace, 1963]. Perhaps as our design techniques become less "empirical" our fundamental ignorance in this area will force us into effort.

3.6 Decision Theory

This is intended to include classical statistical decision theory [Wald, 1950] as well as theories of heuristics and induction [Polya, 1954] [Reichenbach, 1949]. The overlap with what has been described as "Optimization Theory" is obvious, and perhaps the two should appear as one heading. Some good recent attempts to relate decision theory with pattern recognition and neuromime networks can be found in [Tou and Wilcox, 1964]. Many references are given by [Minsky, 1963a].

4. RECAPITULATION OF RECOMMENDATIONS

Throughout this study the question has been borne in mind: "Are the important considerations of artificial intelligence ready for immediate applications in the space sciences?" It appears this question can be answered with an unequivocal "no". Our fundamental ignorance is considerable, and our efforts, to date, mere scratchings compared to what lies ahead. If NASA is to anticipate true humanoid autonomous systems, much more fundamental work must be started and supported. The attempt to build very complex machines to imitate the functions of the human brain, when we do not properly understand either, may be compared to hoisting oneself by one's own bootlaces. The previous sections have indicated the areas where work has been started but where much effort remains to be expended. These will now be briefly recapitulated.

4.1 Replacing the human in the control loop [Fogel, 1963]

Electronic control systems have shown short response times, but have lacked computational sophistication. We must learn how to give these greater ability for decision-making and adaptation. They must be able to predict future system states and operate on "raw" sense data. In spite of this anticipated greatly increased sophistication, these systems should be able to adapt to reasonable degrees of failure of computing or sensing elements, by appropriate "redundancy". [Lee, et al., 1963a] This is particularly important on long missions when maintenance may be a problem.

4.2 Learning-machine research [MacKay, 1961]

A learning-machine can be used for computation or control if it can be trained to converge to some terminal behavior. [Tou and Wilcox, 1964] Some interesting theoretical questions are: "Given a machine, how quickly will it converge?" "What is an optimum training environment?" "How stable is the final result?" "Is further adaptation still possible?" So far, these questions have only been answered, to any extent, for spatial event pattern recognizers (perceptrons, etc.). The really interesting cases, where elements with memory are transducing the reinforcement channels of other such elements, have barely been touched. [Lee, Gilstrap, et al., 1962-63] [Pedelty, 1964, pp. 250-51] Probabilistic machines, dealt with in the first reference cited, need much more attention.

The technology of these machines is lagging considerably and more attention to fundamental physical processes suitable for adaptive machines is required [Pask, 1962] [Pulvari, 1963] [Lee, et al., 1963b]

Such machines can be trained in an artificial environment; or they may be designed to mimic the human operator during an actual on-line control problem. As an interim research step of great value, attention must be given to having better automatic recording and analyzing of these human tasks, performed by "conventional" computing equipment.

In the very near future, further research involving the coupling of learning machines to real environments will be required. These will be the forerunners of "moon-crawlers", etc. These research vehicles would prove valuable for testing theories of adaptation and recognition, particularly where only general criteria regarding "mission" objectives have been supplied (via appropriate goal-hierarchies, in which parts of the machine bias the adaptation of other parts. [Pedelty, 1963b]) A good way of taking advantage of a priori information (for a given class of environments) must be found, and control tasks must be appropriately partitioned to improve learning speed and keep equipment demands within bounds. The work on "systems description" should assist. Better designs may come about when better "man-machine design systems" can be utilized. [Pedelty, 1965]

4.3 Automata Theory

The topics just discussed require better theoretical underpinning from Automata Theory. This subject must concern itself more with the finite adaptive system communicating with a real dynamic environment. This work could not only lead to better systems languages but may provide the key to the difficult problem of effectively utilizing microelectronic techniques.

4.4 Information Theory

More work is required on reliable computation in the presence of noise and/or irregular maintenance cycles.

In the light of a machine's current goals, it is necessary to have measures of the meaning (value) of a given datum. (Strategies can only be properly modified if meaningful data is noted and organized.)

4.5 Systems Theory

Work is just being started on selecting more meaningful indices of performance. It should no longer be necessary to select these to be computationally convenient for the designer, now that better man-machine design systems are just over the "horizon". Obviously, statistics must play a role here.

4.6 Conclusion

The recapitulation section of this report is a condensation of a thumbnail sketch. It can best be seen in proper perspective only in conjunction with the body of the report and the references cited throughout. If this has been the first excursion for the reader in the area of artificial intelligence, he would hardly be blamed if he felt a little bemused. It can only be pointed out that perspective is to be gained only by several years of study of the subject's many facets (some of which, notably the biological and psychological, have barely been touched upon in this report). Those of us who consider we have appropriate backgrounds must inevitably be specialists in only one area, and we can take little comfort from the sea of shifting sand on which we set our course. But progress means change, and we buy this with our educated gambles. Who can say that we shall not soon discover that the earth does not revolve around the sun? This would be painful, but the price of knowledge has always been high, and we shall abandon more than one idea before we can measure our advance.

5. BIBLIOGRAPHY

Arbib, M. A., 1964, Brains, Machines, and Mathematics, McGraw-Hill.

Bernhard, R., 1962, "New considerations in cybernetics research," Cybernetica, 5, No. 1, pp. 5-24.

Chomsky, A. N., 1959, "A note on phrase structure grammars," Information and Control, Dec., 2.

Cowan, J. D., 1962, "Many-valued logics and reliable automata," in [E].

Dimsdale, B., and Markowitz, H., 1964, "A description of the SIMSCRIPT language," IBM Sys. J., 3, No. 1, pp. 57-67.

Fogel, L. J., 1963, Biotechnology: Concepts and Applications, Prentice-Hall.

Friedman, J., 1962, "A decision procedure for computations of finite automata," J.A.C.M., 9, No. 3, pp. 315-323.

Gilstrap, Jr., L. O., and Pedelty, M. J., 1963, "The Epistemological Foundations of Machine Intelligence," report to A.F. Systems Command (Avionics Lab., Dayton, Ohio), ASD-TDR-63-877.

Gunther, G., 1962, "Cybernetic ontology and transjunctional operations" in [G].

Hartmanis, J., 1961, "Task simplification and learning devices," in [I].

Holland, J. H., 1962, "Outline for a Logical Theory of Adaptive Systems," J.A.C.M., 9, No. 3, July 1962.

Holland, J. H., 1964, "Universal Embedding Spaces for Automata," unpublished notes for the Univ. of Michigan Engineering Summer Conf. "Automata Theory," Summer 1964.

Hunt, E. B., and Hovland, C. I., 1963, "Programming a model of human concept formulation," in [D] (see also p. 495).

Krohn, K. B., and Rhodes, J. L., 1963, "Algebraic theory of machines" in [F].

Lee, R. J., Gilstrap, L. O., _et al_., 1962-63, _Theory of Probability State Variable Systems_, Vols. I through VI, ASTIA AD's: 427-872, 428-099, 428-087, 427-775, 428-015, 427-771.

Lee, R. J., et al., 1963a, _Functional Capability of Neuromime Networks for Use in Attitude Stabilization Systems_, A.F. Systems Command Tech. Doc. Report ASD-TDR-63-549, Dayton, Ohio, September.

Lee, R. J., _et al_., 1963b, _Research on Adaptive Sandwiches and Learning Automata_, Adaptronics, Inc., Final Report, Contr. AF 33(657)-8674, Aeronautical Systems Div., May.

Lettvin, J. Y., 1964, Personal communication to the author.

Luisi, J. A., 1962, "Pattern Recognition Applied to Learning in Control Systems," CISL Memo. 62-15, School of E.E., Purdue Univ., Aug. 1962.

MacKay, D. M., 1956, "The epistemological problem for automata," in [A].

MacKay, D. M., 1961, "Information and learning," in [H].

Minsky, M., 1963a, "A selected descriptor-indexed bibliography to the literature on artificial intelligence," in [D].

Minsky, M., 1963b, "Steps toward artificial intelligence," in [D].

Pask, G., 1962, _An Approach to Cybernetics_, New York, Harper and Row.

Pedelty, M. J., 1963a, "A memory complex for an adaptive control system," in [C].

Pedelty, M. J., 1963b, An Approach to Machine Intelligence, Washington, D. C., Spartan Books.

Pedelty, M. J., 1964, "Temporally Organized Automata and an Algebraic Theory of Machines," in [B].

Pedelty, M. J., 1965, Studies of Possibilities of Machines Showing Adaptive Behaviour, as Aids in Computational Tasks, Ph.D. Thesis due to be submitted to the University of London, Fall, 1965 (copies should become available from the author).

Polya, G., 1954, Patterns of Plausible Inference, Princeton Univ. Press.

Pulvari, C. F., 1963, "Ferrielectrics and their application in solid-state devices as an adaptive control," in [J].

Ray-Chaudhuri, D. K., 1961, "On the construction of minimally redundant reliable system designs," Bell Syst. Tech. J., XL, No. 2, Mar., 1961

Reichenbach, H., 1949, The Theory of Probability, Berkeley, Calif., Univ. of Calif. Press.

Sarture, C. W., and Aseltine, J. A., "Performance criteria in adaptive control," in [C].

Simmons, R. F., 1965, "Answering English questions by computer: a survey," Comm. A.C.M., 8, No. 1, January, p. 53 et seq.

Sprague III, C. F., 1964, "On the reticulation problem in multivariable control systems," IEEE Trans. on Automatic Control, AC-9, No. 4, October.

Sutherland, I. E., 1963, "Sketchpad: A man-machine graphical communication system," AFIPS Spring Joint Comp. Conf., Detroit, May, 1963, Spartan Books.

Tou, J., and Wilcox, R., 1964, see [B].

Wald, A., 1950, Statistical Decision Functions, Wiley.

Wallace, A., 1963, Epistemological Foundations of Machine Intelligence, Report to A. F. Systems Command (Avionics Lab., Dayton, Ohio), ASD-TDR-63-878.

Westcott, J. H., 1962, An Exposition of Adaptive Control, p. 3, London, Pergamon Press.

Widrow, B., 1963, "A statistical theory of adaptation," in [C].

Wymore, A. W., 1964, Univ. of Arizona, "The Mathematical Foundations of System Theory," paper presented at 20th National Electronics Conf., Chicago, Ill.

* * * *

[A] Shannon, C. E., and McCarthy, J., (eds.), 1956, Automata Studies, Ann. Math., 34, Princeton Univ. Press.

[B] Tou, J. T., and Wilcox, R. H., (eds.), 1964, Computer and Information Sciences, Washington, D. C., Spartan Books.

[C] Caruthers, F. P., and Levenstein, H., (eds.), 1963, Adaptive Control Systems, MacMillan.

[D] Feigenbaum, E. A., and Feldman, J., (eds.), 1963, Computers and Thought, McGraw-Hill.

[E] von Foerster, H., and Zopf, G., (eds.), 1962, Principles of Self-Organization, Pergamon.

[F] Fox, J., (ed.), 1963, Mathematical Theory of Automata, Microwave Res. Inst., Vol. 12, Brooklyn, N. Y., Polytechnic Press, and Interscience.

[G] Yovitts, M., Jacobi, G. T., and Goldstein, G. D., (eds.), 1962, Self-Organizing Systems, 1962, Spartan Books.

[H] Billings, H., (ed.), 1961, Lernende Automaten, Munich, Oldenbourg.

[I] Cherry, C., (ed.), 1961, _Proceedings of the Fourth London Symposium on Information Theory_, London, Butterworth; also New York, Academic.

[J] Butsch, Jr., L. M., and Oestreicher, H., (eds.), 1963, Special Issue on Bionics, IEEE Trans. on Military Electronics, MIL-7, Nos. 2 & 3, April-July, pp. 84-272.

END