Time Is - Time Was - Time Is Past
Computers for Intelligence

BY HOWARD H. CAMPAIGNE

Unclassified

The "Intelligence" of the sub-title does not mean "military information," which it could very well at a conference* like this, but means rather "adaptive behavior," or "imagination," or "pattern recognition." Frankly, I do not have a single definition for what I mean, but a recurring idea (See Bibliography) has been that some day a machine might be made which exhibited intelligence. Roger Bacon was the first to succeed, it was said, but his machine refused to discuss trivialities with man and then destroyed itself in frustration at man's inability to communicate with it. The story of Roger Bacon is that he succeeded in building an artificial intelligence, probably in the year 1277 or just before that, for that was the year in which he was arrested and imprisoned, charged with "innovations." It was a defense project, the ultimate objective being to build a wall of brass around England. His intelligence was housed in an artificial head. It took him 7 years to build it—about right for a defense project. When it was done, he spent 60 days debugging, mostly overtime. That sounds familiar. This debugging stretch was ended by the irresistible need for sleep, so he left his assistant in charge. That clown could hardly wait for Friar Bacon to fall asleep in order to push the start button. The head said "Time is," and lit the halt light. Some clownish talk and another push on the button elicited "Time was." More irrelevant comments and a button push caused the shout "Time is past," and the machine smashed itself on the floor. That was the end of the project. Clearly the head was about to say that to get support they needed the term a "real time system," but it could not make itself understood. This scene of acute frustration has typified artificial intelligence ever since.

The report on this was written by Robert Greene in 1588, 311 years later, almost a record delay for a progress report. It was typed in 1592 and released in 1630. A bad precedent.

If we are to have a demonstration of intelligence by a machine, we must agree in advance on what constiutes an adequate demonstration. I have talked to some of my colleagues about this, and I despair of ever getting any agreement. By and large there is some

* Read before the MIL–E–CON 8, 16 Sept 1964.

consensus that by an intelligent machine we mean one which exhibits human behavior of some kind. Turing[1] reduced this to a contest with teleprinters (a weapon which computers can use readily) and defined intelligence as the ability to imitate a man. Kelly and Selfridge[2] suggested an even simpler game. But beyond this there is no agreement. Some of my friends would be satisfied that a machine was intelligent if it could outperform some human being. But there are human beings–present company excepted–whose performance is miserably low, and this standard may already have been met. Others, also my friends, would not admit intelligence in a machine unless it excelled all human beings. If it could do this, then who would be the judge?

If we are to demonstrate intelligence in a machine, we must decide what we mean. To do this we can start at either end; what can a machine do or what can a man do? The latter is not nearly as well understood as one would think. The abilities lumped together under the word "intelligence" are various and not ordinarily distinguished clearly one from another. Occasionally an "idiot savant" appears who demonstrates vividly that intelligence is composite. He can do arithmetic with great ease but is unable to comprehend social relations, or he has great skill in music but none in language, or exhibits some other such striking disparity in abilities. Finding what these abilities are is the unfinished job of the psychologist.

These two approaches are like digging a tunnel by starting at both ends, although in this case we know so little of the mountain that we don't know that these diggers are even going toward each other, let alone whether they will meet. And if, by accident, they should meet, we don't know of what use the tunnel will be!

To my mind, the more productive way is to start with the machine and find what limits its ability. Even if such a program has a negative result and shows that intelligence is not achievable by a machine and that man is able to do something of a higher type, it will be useful to know the boundaries. I do not expect this to be the result. I think that as we learn more of what machines can do and more about what is rational behavior by human beings the question will go away.

A useful analogy is the development of flying. For centuries men dreamed of imitating birds. DaVinci made drawings of linkages which would work a wing DaVinci wrote "A bird is an instrument working according to Mathematical law, which instrument it is

[1] Alan Turing, "Computing Machinery and Intelligence" *Mind*, 59, New Series 236, October 1950, pp. 433–460.

[2] J. L. Kelly, Jr., and O. G. Selfridge, "Sophistication in computers: a disagreement," *IRE Transactions on Information Theory*, 1962, IT–8, pp. 78–80

within the capacity of men to reproduce with all its movements." At that time arguments were advanced that "birds are supported by the hand of God" and "if He meant us to fly, God would have given us wings." In the end it was components unknown in nature, the propeller and the internal combustion engine, that opened the way to a solution. These components were developed by people more interested in what could be done with engines than they were in imitating birds. The Wright brothers' solution was not satisfactory to all concerned, and interest in ornithopters continued for a while, but now the question of how closely a man can imitate a bird is dormant.

Meszar[3] said "any mental process which can be adequately reproduced by automatic systems is not thinking." Since then we have had the following demonstrations: theorem proving by Gelernter, checkers playing by Samuel, music composition by Hiller and Isaacson, assembly line balancing by Tonge in 1961, designing motors by Gold in 1959, and freshman calculus by Slagle in 1961.

Minsky[4] has listed, as clearly as the muddled state of the art permits, operations which must be performable by any thinking machine. He describes these as search, pattern recognition, learning, planning, and induction. The order cited is that of increasing sophistication, that is to say, of decreasing understanding. There have been respectable demonstrations of each of these except the last, induction. Search has been implemented and written about by a great number of investigators; the simplest of the concepts, it still furnishes much discussion and is not to be disposed of in the near future.[5] Pattern recognition here means matching against a prototype; this is being done commercially as in the reading of checks written with magnetic ink, although the limitations on the technique are not well understood. Learning is here restricted to adaptive behavior, which in its simplest form can be easily demonstrated, but which titilates us in our limited ability to generalize; in its most general form it would solve all of our problems. If we equate adaptive behavior with learning and if we assume no limit on learning (why should it be limited?), then, as I. J. Good says, this is the last invention man will ever have to make. What Minsky calls "planning," I would call "reformulation." Demonstrations of this have been made by Newell, Shaw and Simon,[6]

[3] J. Meszar, "Switching Systems as Mechanized Brains," *Bell Telephone Record*, February 1953.

[4] Marvin Minsky, "Steps Toward Artificial Intelligence," *Proc IRE*, 49, 1961, pp. 8–30.

[5] For example, C. E. Shannon, "Programming a Computer to Play Chess," Philosophical Magazine, 7, 41, 1950, pp. 225-275.

[6] Alan Newell, J. C. Shaw and H. A. Simon, "Empirical explorations of the logic theory machine," *Proc WJCC*, 1957, pp. 218-230.

who show a way of finding pertinent intermediate goals. The last in his list is induction; this is the $64 question as we see it today; no demonstrations have been made that I know of. Another name for induction is jumping to conclusions.

Newell says "if ever a fully capable intelligent program is realized, it will be recognized by noting that it can get along without any programmer at all."[7]

Minsky says "Almost any problem can be converted into a problem of finding a chain between two terminal expressions in some formal system."[8]

Let me go over Minsky's list again, this time looking toward the directions in which these techniques, if we had them, would generalize. The problem of search, listed by Minsky as the first step in his program toward artificial intelligence, has great generality and has had many contributions from many sources. Hill climbing, the method of steepest descent, heuristic methods, all are attempts to find ways to exploit the structure of the space being searched. Basic as this technique is, it still may be of complete generality, for any solution can be stated in terms of a choice of paths. If one thinks of search as analogous to a man exploring territory, then the amount he can see at one time is the important parameter. Crossing the plains toward a distinctive peak in daylight is different from thrashing through the jungle in the dark.

The term "pattern recognition" could be interpreted to mean the perception of similarities in previously undigested data. In this broad sense it is very much like, perhaps equivalent to, induction. It was not used in this sense by Minsky. In the narrower sense of recognizing a resemblance to a prototype it still is powerful in categorizing concepts.

"Learning" too is often used more broadly than Minsky used it. A device which was adaptive in the broadest sense could accommodate to any situation, barring a catastrophe. The biological ecological system of evolution may be doing just this. The analogy between evolution and learning is a striking one but also painful because of the slow reaction of the first. Evolution is a blind search. The ecological system is searching in parallel, of course, but each species is trying to solve its own problems alone. If one thinks of the genetic possibilities as a space with as many dimensions as there are genes and

[7] M. C. Yovits, G. T. Jacobi, G. D. Goldstein, "Some problems of basic organization in problem-solving programs," *Self-organizing Systems*, Spartan Books, Baltimore, 1962.

[8] Marvin Minsky "Steps Toward Artificial Intelligence," *Proc IRE*, 49, 1961, pp. 8–30.

as many points on each axis as there are alleles, then each individual's heritage can be represented by a point in this space. In the case of a monosexual species (if there are any), each strain is making its own path. Bisexual species advance in a herd, each specimen being at a point in the genetic space A herd can split, of course, but the splitting is limited by the fact that each species is at a local maximum, or trap. When dislodged from its cul-de-sac, a species mutates fairly rapidly and generally toward what seems like conscious direction, uphill perhaps, a phenomenon noticed long ago by geneticists and called "orthogenesis."

By "planning" Minsky means the substitution of intermediate goals as means to an ultimate goal. This is akin to reformulation, if not the same, a step which I am sure is essential to human problem solving. Our experience in scientific research is that once a question has been properly put, the answer generally follows in short order. A researcher spends more effort wrestling with questions than he does with answers. Those questions which seem to be well put but for which no answer has yet been found are famous, such as the four colored map problem. This substitution of goals has been demonstrated by Newell, Simon and Shaw[9]. In the context of searching through a graph, it is easily seen that the intermediate goal is an island of great value. If one had a problem which required the determination of six parameters, each with forty possible values, then a blind search would be faced with four billion places to look. If a sub-goal can be found so that three parameters can be determined first and then the other three, we find that only $2 \times 64{,}000$ places need be searched, or 128,000. The existence of the sub-goal is worth a factor of 32,000. The existence of such goals depends on the structure of the problem, of course.

Induction is the reasoning from a part to the whole, the predicting of new events from past events. To do this predicting, one needs a model of the world, or of the relevant parts of the world. The construction of the mathematical model from a sample is the step we do not understand. It is jumping to a conclusion, which, up to now, machines do not do so well. The most formal kind of induction, mathematical induction, is a special case of deduction, an operation done very well and commonly by machines. Perhaps there is little difference between the two kinds of reasoning, and our fears of the unknown are not justified.

When I started this paper, I meant to work toward specifying a machine or computer useful in experiments on artificial intelligence. If we are going to experiment with thinking computers, what kind of

[9] Newell, Shaw and Simon, *op. cit.*

machines would we like? Software improvements are a necessity, but a number of other things suggest themselves.

One is speed. If our computer could explore our game of chess to the very end, then, of course, it would have insight; but this is impossible with chess and with many other problems, because of the tremendously ramified argument. Chess has been quoted as having 10^{120} positions, and if only a millionth of these were relevant and if they could be disposed of at a million a second, it would take 3×10^{100} years to exhaust them. So speed by itself will not do much; it will take software.

Another is parallelism. Selfridge[10] has described an organization resembling a situation room which has a big board on which the latest data is available to all, and which can be continually updated by each of a large number of demons: pandemonium. Such an organization might have advantages in learning. There is reason to think that the human brain may be a little like this, a committee of slow components. The programming of such an organization is almost unexplored.

An alternative is that of distributed logic, an assemblage of data-manipulating equipment, especially memory, each of which can do some of the essential processes such as sensing and combining Thus the logic, instead of being concentrated in the accumulator or control, would be everywhere. A content addressed memory is a kind of distributed logic; with this, one can do many complex operations, such as sorting, almost painlessly. Improvements in software have more to offer.

But a useful thinking machine must have flexible input-output, an effective interface with men. Like Bacon's head, if it cannot get through to us, it might as well not exist. And we too have our language problems and need to have the very best of aids in stating problems to the device and reorganizing the thinking of our machine. This is the area which needs most improvement, the easy interchange of information between man and machine.

Buchman also has characterized the various properties of an intelligent machine in a different but very clear way.[11] He says that such a machine must be adaptive, self-organizing, or learning. By "adaptive" he means stable in a changing environment; by "self-organizing" he means effective in a radically changing environment;

[10] Oliver Selfridge, "Pandemonium: A Paradigm for Learning," *Paper 3-4*, Teddington Symposium, November 1958.

[11] A. F. Buchman, "The Digital Computer in a Real-Time Control System, Part III," *Computer Design*, Vol. III, No. 5, May 1954, pp. 24-31.

and by "learning" he means increasingly effective in a stable environment.

I must comment on a statement I have seen that soon the chess champion will be a machine! This is fatuous. Bicycles are not used in the Olympic footraces; if they were, a cyclist would be world champion. When the rules of chess are amended to prohibit mechanical aids, that will be a clue that one of our subgoals is being approached.

May I suggest that Turing's test with the game of bridge might be effective? Played by teletype, it would be the task of each player to identify the machine among the three other players. Or if bridge is too much work to program, then a series of checkers games, where a man plays alternately with a champion and with Samuel's superlative checker program, the man's task being to name which is which. This could be implemented readily because Samuel's part is done.

If a thinking machine can be built, then it must be done; it is a matter of self-respect. Just as a man must be put on the moon, just as Mount Everest had to be climbed, just as the poles had to be visited, just as a flying machine had to be made no matter what the arguments against it, so a machine must be made which can think.

Taube has said " . . . The proper man-machine relation is one of complementation . . . "[12] I do not gain-say this; I agree. But the demonstration must be made nevertheless. Seashore's story illustrates the state of the art.

[12] Mortimer Taube, *Computers and Common Sense: The Myth of Thinking Machines*, Columbia University Press, 1961.

BIBLIOGRAPHY ON EXTENDING SCOPE OF COMPUTERS

S. Amarel, "On the Automatic Formation of a Computer Program which Represents a Theory," *Self Organizing Systems*, Spartan Books, Washington, D. C., 1962.

A. M. Andrew, "Learning Machines," Paper 3-6, Symposium on the Mechanization of Thought Processes, Teddington, England, November 1958.

James B. Angell, *The Need and Means for Self-Repairing Circuits*, Technical Report No. 4654-2, USAF Contract AF33(616)-7726, Stanford Electronics Laboratories.

Paul Armer, "Attitudes toward Intelligent Machines," *Datamation*, Vol. 9, No. 3, March 1963, pp. 34-38.

————, "Attitudes toward Intelligent Machines," RAND Corp., p 2114, 30 September 1960. (Extensive bibliography.)

J. A. Aseltine, A. R. Mancini, and C. W. Sarture, "Impulse-Response Self-Optimization as Compared with Other Criteria for Adaptive Systems," Presented at 4th Annual Instruments and Regulators Conference on Automatic Optimization, University of Delaware, 4 April 1958.

W. R. Ashby, *Design for a Brain*, John Wiley & Sons, Inc., New York, N. Y., 1952. (Chapman and Hall, London.)

————, "Design for a Brain," *Electronics Engineering*, 20, 1948, pp. 379-383.

————, "Computers and Decision Making," *New Scientist*, 7, 746, 24 March 1960.

————, "Design for an Intelligence Amplifier," *Automata Studies, Annals of Mathematical Studies*, No. 34, Princeton University Press, 1956, p. 215.

————, "What is an Intelligent Machine," *Proc. WJCC*, Vol. 19, May 1961, p. 275.

————, and J. Rignet, "The Avoidance of Over-Writing in Self-Organizing Systems," Technical Report No. 1, Burden Neurological Institute, Bristol, England.

M. L. Babcock *et al*, "Some Principles of Preorganization in Self-Organizing Systems," Electrical Engineering Research Laboratory Report No. 2, University of Illinois, 24 June 1960.

R. B. Banerji, "Computer Programs for the Generation of New Concepts from Old Ones," Case Institute of Technology. (Preprint)

H. B. Barlow, "Sensory Mechanisms, the Reduction of Redundancy, and Intelligence," Symposium on the Mechanization of Thought Processes, Teddington, England, November 1958.

B. L. Basore and W. D. Wood, *A Model for Communication with Learning*, Dikewood Corporation TN-1004-2, 31 May 1960.

A. Bernstein *et al*, "A Chess Playing Program for the IBM 704," *Proc. WJCC*, 1958, pp. 157-159.

A. Bernstein and M. deV. Roberts, "Computer vs. Chess Player," *Scientific American*, 198, June 1958, pp. 96-98.

"Bibliography on Biological and Artificial Intelligence," Jet Propulsion Laboratory, California Technical Literature Search No. 254 and Supplement.

W. W. Bledsoe and I. Browning, "Pattern Recognition and Reading By Machine," *EJCC*, Boston, December 1959, pp. 225 232.

R. R. Bush and F. Mosteller, *Stochastic Models for Learning*, John Wiley & Sons, Inc., New York, N. Y., 1955.

Silvio Ceccato, "La Machine qui Parle et qui Pense," *Congress International de Cybernetique*, Namur, 1956.

C. K. Chow, "An Optimum Character Recognition System using Decision Functions," *IRE Trans. on Electronic Computers*, Vol. EC6, December 1957, pp. 247–254.

W. A. Clark and B. G. Farley, "Generalization of Pattern Recognition in a Self-Organizing System," *Proc. WJCC*, 1955, pp. 86–91.

W. E. Dickinson, "A Character-Recognition Study," *IBM Research and Development*, Vol. 4, July 1960, pp. 335–348.

T. L. Dimond, "Devices for Reading Handwritten Characters," *Proc. EJCC*, Washington, D. C., December 1957, pp. 232–237.

G. P. Dineen, "Programming Pattern Recognition," *Proc. WJCC*, March 1955, pp. 94–100.

W. Doyle, "Recognition of Sloppy Hand-printed Characters," Lincoln Laboratory Group Report 54–12, December 1959.

T. G. Evans, *A Heuristic Program for Solving Geometric Analogy Problems*, unpublished Doctoral Disseration, MIT, 1963. (Also Spring *JCC*, 1964, pp. 327–338.)

B. G. Farley and W. A. Clarke, "Simulation of Self-Organizing Systems by Digital Computers," *Transactions on Information Theory, IRE PGIT*, 4, September 1954, pp. 76–84.

Edward A. Feigenbaum, "Artificial Intelligence Research," *IEEE Transactions*, Vol. IT–9, No. 4, October 1963, pp. 248–253.

E. Feigenbaum and J. Fedman, Eds., *Computers and Thought*, McGraw-Hill, N. Y., 1963.

Edward A. Feigenbaum and Gerbert A. Simon, "Forgetting in an Association Memory," *RAND* p–2311, 24 May 1961.

I. Flore and L. Grey, "Optimization of Reference Signals for Character Recognition Systems," *IRE Transactions on Electronic Computers*, Vol. EC–9, March 1960, pp. 54–61.

L. J. Fogel, "Toward Inductive Inference Automata," *Proc. 1962 International Conference of Information Processing*, Amsterdam, The Netherlands, 1963, pp. 395–400.

R. M. Friedberg, "A Learning Machine: Part I," *IBM Journal of R&D*, 2, No. 1, January 1958, pp. 2–13.

R. M. Friedberg, B. Dunham and J. H. North, "A Learning Machine: Part II" *IBM Journal of R&D*, 3, July 1959, pp. 282–287.

H. Gelernter and N. Rochester, "Intelligent Behavior in Problem-Solving Machines," *IBM Journal of R&D*, 2, 1958, pp. 336–345.

H. L. Gelernter, "Realization of a Geometry Theorem-Proving Machine," *Proc. Int. Conf. on Information Processing*, Paris, 1959.

———, "Theorem Proving by Machine," *IBM IR–00124*, August 1957.

A. Gill, "Minimum-scan Pattern Recognition," *IRE Transactions on Information Theory*, Vol. IT–5, June 1959, pp. 52–58.

———, "Possibilities for the Practical Utilization of Learning Processes," Paper 4–10, Symposium on Mechanization of Thought Processes, Teddington, November 1958.

P. C. Gilmore, "A program for the production of proofs for theorems derivable within the first order predicate calculus from axioms," Unesco, NS, ICIP, 1.6.14, International Conf. on Information Processing, Paris, June 1959.

———, "A Proof Method for Quantification Theory," *IBM Journal of R&D*, 4, 1960, pp. 28–35.

H. T. Glantz, "On the Recognition of Information with a Digital Computer," *JACM*, Vol. 4, April 1957, pp. 178–189.

B. Gold, "Machine Recognition of Hand-sent Morse Code," *IRE Transactions on Information Theory*, Vol. IT–5, March 1959, pp. 17–24.

I. J. Good, "The Subassembly Theory of Memory and Meaning and its Relevance to the Economical Construction of an Ultra-Intelligent Machine," May 1963.

S. Gom, "On the Mechanical Simulation of Learning and Habit Forming," *Information and Control*, 2, September 1959, pp. 226–259.

B. F. Green, A. Wolf, C. Chomsky, and K. Langhevy, "Baseball: an Automatic Question Answerer," *Proc. WJCC* 1961, pp. 219–224.

P. H. Greene, "An Approach to Computers that Perceive, Learn and Reason," *Proc. WJCC*, 1959, pp. 181–186.

———, "A Suggested Model for Information Representation in a Computer that Perceives, Learns and Reasons," *Proc. WJCC*, 1960, pp. 151–164.

———, "Networks for Pattern Perception," *Proc. National Electronics Conf.*, Vol. 15, October 1959, pp. 357–369.

E. C. Greenian and Y. M. Hill, "Considerations in the Design of Character Recognition Devices," *IRE National Convention Record*, 1957, pp. 119–126.

E. C. Greenian *et al*, "Design of Logic for Recognition of Printed Characters by Simulation," *IBM Journal of R&D*, Vol 1, January 1957, pp. 8–18.

R. L. Grimsdale *et al*, "A System for the Automatic Recognition of Patterns," *Proc. IEE*, 106, Pt. B, March 1959.

Fred Gruneberger, "Benchmarks in Artificial Intelligence," *Datamation*, October 1962, pp. 33–35.

George L. Haller, "Our State of Mind in 2012 A. D." *Proc. of IRE*, Vol. 50, No. 5, May 1962, pp. 624–627.

R. W. Hamming, "Intellectual Implication of the Computer Revolution," *Bell Telephone Laboratories*.

L. D. Hamon, "A Line-drawing Pattern-recognizer" *Proc. WJCC*, San Francisco, Calif., May 1960, pp. 351–364.

C. C. Heasly, "Some Communication Aspects of Character-sensing Systems," *Proc. WJCC*, San Francisco, Calif., May 1959, pp. 176–180.

W. H. Highleyman and L. A. Kamentsky, "Comments on a Character Recognition Method of Bledsoe and Browning," *IRE Transactions on Electronic Computers*, EC9, June 1960, p. 263.

M. E. Hoff, Jr., "Learning Phenomena in Networks of Adaptive Switching Circuits," Technical Report No. SEL–62–090, Stanford Electronics Labs., July 1962.

Aiko M. Hormann, "Programs for Machine Learning," TM–669/000/01 SDC 29 May 1962.

E. B. Hunt, *Concept Formation: An Information Processing Problem*, John Wiley & Sons, Inc., New York, 1962.

F. H. Jean, *Generation and Testing of Hypotheses*, Dikewood Corporation, FR–1021, Contract AF30(602)–2514 for Rome Air Development Center, 29 May 1962.

L. A. Kamentsky, "Pattern and Character Recognition Systems—Picture Processing by Nets of Neuron-like Elements," *Proc. WJCC*, San Francisco, May 1959, pp. 304–309.

J. L. Kelly, Jr., and O. G. Selfridge, "Sophistication in Computers: a Disagreement," *IRE Transactions on Information Theory*, IT–8, 1962, pp. 78–80.

T. Killren, R. L. Grimsdale, and F. H. Summer, *Experiments in Machine Learning and Thinking*, UNESCO/NS/1C1PJ.6.15, 1959.

R. A. Kirsch, C. Ray, L. Cahn, and G. H. Urban, "Experiments in Processing Pictorial Information with a Digital Computer," *Proc. EJCC, Proc. IRE.*, December 1957, pp. 221–229.

J. Kister, P. Stein, S. Ulam, W. Welden and M. Wells, "Experiments in Chess," *JACM.*, April 1957.

S. Kuroda, "An Investigation on the Logical Structure of Mathematics (XIII)—A Method of Programming of Proofs in Mathematics for Electronic Computing Machines," *Nagoya Mathematical Journal*, 16, 1960, pp. 145–203.

P. A. Lachenbruch, A. J. Slevenske and A. C. Marchese, "Artificial Intelligence—A Summary of Current Research & Development," American Institute for Research, Los Angeles, No. AIR–C63–2/62–TR, 1962.

J. C. R. Licklider, "Interactions Between Artificial Intelligence, Military Intelligence and Command and Control," preprints 1st Congress of the Information System Sciences, MITRE Corp., Bedford, Mass., 1962.

D. M. MacKay, "Mindlike Behaviour in Artefacts," *British Journal for one Philosophy of Science*, II, 1951, pp. 105–121.

Carl Maltz, "A Measure of the Significance of Pattern Features for Use as an Aid in the Design of Recognition Systems," Report 62–68, University of California, Los Angeles, December 1962.

T. Marrill and D. M. Green, "Statistical Recognition Functions and the Design of Pattern Recognizers," *IRE Transactions on Electronic Computers*, Vol. EC–9, No. 4, December 1960, pp. 472–477.

Henrik H. Martens, "Two Notes on Machine 'Learning,'" *Information and Control*, 2, 1959, pp. 364–379.

R. L. Mattson, "A Self-organizing Binary System," *Proc. EJCC*, 1959, pp. 212–217.

J. McCarthy, "Programs with Common Sense," *Mechanisation of Thought Processes*, Vol. I, National Physical Laboratory Symposium No. 10, Her Majesty's Stationery Office, 1959, pp. 75–84.

Warren S. McCuloch, "The Brain as a Computing Machine," *Electronics Engineering*, 69, 1949, p. 492.

J. Mesyar, "Switching Systems as Mechanized Brains," *Bell Telephone Laboratories Record*, February 1953.

Margaret Milligan, "Machines are Smarter Than I Am!" *Data Processing Digest*, October 1959.

O. N. Minot, "Automatic Devices for Recognition of Visible Two-dimensional Patterns; A Survey of the Field," US Naval Electronics Laboratory Technical Memo. 364, San Diego, June 1959.

M. L. Minsky, Appendix to "Steps Toward Artificial Intelligence," preprints, 1st Congress of the Information System Sciences, MITRE, 1962.

————, "A Selected Descriptor-Indexed Bibliography to the Literature on Artificial Intelligence," *IRE Transactions on Human Factors in Electronics*, Vol. HFE–2, No. 1, March 1961, pp. 39–56.

———, "Heuristic Aspects of the Artificial Intelligence Problem," Lincoln Laboratory Report 34–55, December 1956; ASTIA Doc. No. 236885, December 1956.

———, "Learning Systems and Artificial Intelligence." *Applications of Logic to Advanced Digital Computer Programming*, University of Michigan, Coll. of Eng., Summer Session, 1957.

———, "Some Methods of Artificial Intelligence and Heuristic Programming," *Proc. Symposium on the Mechanization of Thought Processes*, NPL Teddington, November 1958.

———, "Steps Toward Artificial Intelligence," *Proc. IRE* 49, January 1961, pp. 8–30.

———, and O. G. Selfridge, "Learning in Random Nets," *Fourth London Symposium on Information Theory*.

E. F. Moore, "On the Shortest Path Through a Maze," *Proc. International Symposium on the Theory of Switching*, Harvard, 1959.

G. A. Morton, "Machines with Imagination," *Proc. IRE*, Vol. 50, No. 5, May 1962, p. 611.

O. H. Mowrer, *Learning Theory and the Symbolic Processes*, John Wiley & Sons, Inc., New York, 1960.

Allen Newell, "Lectures on Heuristic Programs," *Engineering Summer Conferences, Summer Series*, Ann Arbor, Michigan, 1957.

———, "On Programming a Highly Parallel Machine to be an Intelligent Technician," *Proc. WJCC*, Paper 9.3, May 1960, pp. 267–282.

———, "The Chess Machine; An Example of Dealing with a Complex Task by Adaptation," *Proc. WJCC*, March 1955.

——— and H. A. Simon, "A Program that Simulates Human Thought," *Lemende Automaten*, H. Billing, Ed., Oldenbourg, Munich, 1961.

——— and H. A. Simon, "Computer Simulation of Human Thought," *Science*, Vol. 134, 22 December 1961, p. 2011.

———, J. C. Shaw, and H. A. Simon, "A General Problem-solving Program for a Computer," *Computers and Automation*, 8, 1959, pp. 10–17.

———, J. C. Shaw, and H. A. Simon, "A Variety of Intelligent Learning in a General Problem Solver," *Self-Organizing Systems*, M. C. Yovits and S. Cameron (Eds.), Pergamon Press, London, 1960, pp. 153–189.

———, J. C. Shaw, and H. A. Simon, "Chess-Playing Programs and the Problem of Complexity," *IBM Journal of R&D.*, 2, October 1958, pp. 320–335.

———, J. C. Shaw, and H. A. Simon, "Problem Solving in Humans and Computers," *RAND Corp.*, P–987, 7 December 1956.

———, J. C. Shaw, and H. A. Simon, *Report on a General Problem-Solving Program*, UNESCO/NS/1C1P/1.6.8; *Proc. International Conference on Information Processing*, Paris, pp. 256–264.

J. D. North, "The National Behavior of Mechanically Extended Man" Boulton Paul Aircraft Ltd., Wolverhampton, England, September 1954.

A. G. Oettinger, "Programming a Digital Computer to Learn," *Phil Mag.*, Vol. 43, December 1952, pp. 1243–1263.

———, "Simple Learning by a Digital Computer," *IRE Proceedings of the Association for Computing Machinery*, Toronto, Ontario, September 1952.

Proceedings of the Symposium on Mechanization of Thought Processes, H. M. Stationery Office, London, 1959.

W. C. Ridgway III, "An Adaptive Logic System with Generalizing Properties," Technical Report No. SEL–62–040, Stanford Electronics Labs., April, 1962.

P. I. Richards, "On Game Learning Machines," *Scientific Monthly*, 74, 4, 1952, pp. 201–205.

L. G. Roberts, "Pattern Recognition with an Adaptive Network," *IRE International Convention Record*, Pt. 2., 1960, pp. 66–70.

A. Samuel, "Appendix: Game of Checkers Played by Mr. R. W. Nealy vs Samuel Checker Playing Program," *Computers and Thought*, Feigenbaum & Feldman, 1963.

———, "Some Studies in Machine Learning, Using the Game of Checkers," *IBM Journal of R&D*, No. 3, July 1959, pp. 210–229.

O. G. Selfridge, "Pandemonium: A Paradigm for Learning," Papers 3–4, Symposium on the Mechanization of Thought Processes, Teddington, England, November 1958.

———, "Pattern Recognition and Modern Computers," *Proc. WJCC*, March 1955, pp. 91–93.

———, and G. P. Dinneen, "Programming Pattern Recognition," *Proc. WJCC*, March 1955.

———, and U. Neisser, "Pattern Recognition by Machine," *Scientific American*, 203, August 1960, pp. 60–68.

C. E. Shannon, "Game-playing Machines," *Journal of the Franklin Institute*, 206, December 1955, pp. 447–453.

———, "Programming a Computer to Play Chess," *Phil. Mag.* 7, 41, March 1950, pp. 256–275.

H. Sherman, "A Quasi-Topological Method for Recognition of Line Patterns," Unesco, NS, 1C1P, H. L. 5, International Conference on Information Processing, Paris, June 1959.

R. F. Simmons, "Syntex: Toward Computer Synthesis of Human Language Behavior," *Computer Applications in the Behavioral Sciences*, H. Barko Ed., Prentice-Hall, Inc., Englewood Cliffs, N. J., 1962.

Herbert A. Simon, "The Heuristic Compiler," *RAND*, Santa Monica, RM–3588 PR.

———, "Prediction and Hindsight as Confirmatory Evidence," *Phil. of Science*, 22, 1953, pp. 227–230.

———, and Allen Newell, "Computer Simulation of Human Thinking and Problem Solving," *Datamation*, 7 June 1961, pp. 18–20.

N. Sluckin, *Minds and Machines*, Penguin.

R. J. Solomonoff, "An Inductive Interference Machine," *IRE National Convention Record*, Pt. 2, 5, 1957, pp. 56–62.

———, "A Preliminary Report on a General Theory of Inductive Inference," Zator Technical Bulletin, v–131, ZTB–138, February 1960.

———, "The Mechanization of Linguistic Learning," *Zator Technical Bulletin*, No. 125, September 1958.

S. D. Steams, "A Method for the Design of Pattern Recognition Logic," *IRE Transactions on Electronic Computers*, Vol. EC–9, March 1960, pp. 48–53.

Mary Elizabeth Stevens, *Abstract Shape Recognition by Machine*, AFIPS 20, *Proc. EJCC*, 1961, Washington, D. C.

Donald N. Streeter, and Kumpati S. Narenda, *A Self-organizing Control System Based on Correlation Techniques and Selective Reinforcement*, Technical Report

No. 359, Craft Laboratory, Harvard, ONR Contract Nonr1866(16), NR–372–012, 20 July 1962.

D. L. Szekely, "On Basic Aspects of the Concept Transforming Machine," *Cybernetica*, Vol. 4, No. 2, 1961.

M. Taube, *Computers and Common Sense, The Myth of Thinking Machines*, Columbia University Press, 1961.

W. K. Taylor, "Pattern Recognition by Means of Automatic Analog Equipment," *Proc. IEE*, Vol. 106, Pt. B, March 1959.

————, "Electrical Simulation of Some Nervous System Functional Activities," *Information Theory*, C. Cherry Ed., Butterworth Scientific Publications, London, 1956.

W. H. Thorpe, "The Concepts of Learning and their Relation to Those of Instinct," Symposium of the Society for Experimental Biology, IV, p. 387.

J. H. Troll, "The Thinking of Men and Machines," *Atlantic Monthly*, July 1954.

A. M. Turing, "Can a Machine Think," *World of Mathematics*, James R. Newman, Ed., Simon and Schuster, 1956, Vol. 4, p. 2109.

————, "Computing Machinery and Intelligence," *World of Mathematics* Vol. 4, p. 2099; *Mind*, 59, October 1950, pp. 433–460.

L. Uhr, "Latest Methods for Conception and Education of Intelligent Machines," *Behavioral Science*, 4, 1959, pp. 248–251.

————, "Intelligence in Computing Machines; the Psychology of Perception in People and in Machines," *Behavioral Science*, 5, 1960, pp. 177–182.

S. H. Unger, "Pattern Detection and Recognition," *Proc. IRE*, 47, October 1959, pp. 1737–1752.

A. M. Uttley, "Imitation of Pattern Recognition and Trial-and-error Learning in a Conditional Probability Computer," *Rev. Mod. Phys.*, Vol. 31, April 1959, pp. 546–548.

B. Widrow and M. E. Hoff, "Adaptive Switching Circuits," Standard Electronics Laboratory, Stanford, Technical Report No. 1553–1, June 1960.

J. D. Williams, "Toward Intelligent Machines," *RAND Corporation P-2170*, 29 December 1960.

M. T. Yovitts and S. Cameron, *Self-Organizing Systems*, Pergamon Press, New York, 1960.