



# AI'S HALF CENTURY: ON THE THRESHOLDS OF THE DARTMOUTH CONFERENCE

ROBERTO CORDESCHI

Dipartimento di Studi Filosofici ed Epistemologici - Università degli Studi di Roma "La Sapienza"

## 1. Introduction

As is well known, the 1956 summer Dartmouth Conference on AI was preceded by a preparatory document dated August 31, 1955, whose authors were John McCarthy, Marvin Minsky, Nathaniel Rochester and Claude Shannon. The meeting's aim was to examine "the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it", as one reads in the document [9]. Some of the main pioneers in computer programming were present at Dartmouth, such as Allen Newell, Arthur Samuel, Oliver Selfridge and Herbert Simon. After Dartmouth, the historical centres of AI research would be formed: at Carnegie-Mellon University with Newell and Simon, at the MIT with Minsky, and Stanford University with McCarthy. In England, Alan Turing's legacy was taken up by Donald Michie at Edinburgh, before AI research spread to other European countries and around the world.

Differently motivated analyses of the origins and developments of AI have been suggested (see [4], [6], [10], [14]). In the present paper some less well-known contributions and events that precede the Dartmouth Conference are investigated, with the aim of showing how earlier attempts in mechanising intelligence raised several questions which were to become controversial and much debated issues in AI research over the following years.

## 2. Simulating intelligent functions on computers

With the aim of briefly giving the context of my investigation into the origins of AI, let me start with The Computer Issue, a special issue of the Proceedings of the IRE (Institute of Radio Engineers), published with the collaboration of the PGEC (the IRE Professional Group on Electronic Computers) in October 1953. This special issue provides excellent evidence of results in computer design and technology achieved in the 1950s. It included,

among the others, an article by Claude Shannon, "Computers and automata" (a review of computer performances comparable to those of humans: see [17]), and a long series of articles describing digital computers in all their aspects, as regards both software and hardware. In several of these articles there were glimpses of the advantages stemming from the imminent spread of transistors, which, by replacing the cumbersome and unreliable vacuum tubes, would characterise second-generation computers.

The building and dissemination of computers in the United States and Europe was strongly sponsored by government and industry. In the United States, IBM had already supported Howard Aiken's projects in the 1940s. Starting from the 1950s, almost at the same time as Ferranti was completing the Mark 1 computer in England, IBM began producing the type 701 computer, which was carefully described in the Computer Issue. This was the first in a series of electronic general-purpose, stored-program computers which would be used for both theoretical research aims and government and industrial applications. As a researcher at IBM, Nathaniel Rochester, then one of the proponents of the Dartmouth Conference, was responsible for the logical organisation of the type 701, and wrote the first assembly program for it. In 1952, the first checkers program by Arthur Samuel, the author of the opening article for the Computer Issue, was run on this computer.

This and other programs were illustrated by Shannon in his article in the Computer Issue, including the checkers program by Christopher Strachey, who had published a report in 1952. Other programs were able to play games fairly well: the program by D.W. Davies for tic-tac-toe, which ran on a DEUCE computer, and that for nim, running on the NIMROD electronic computer, built by Ferranti. In 1954, Samuel completed the implementation of the first learning checkers program on an IBM 704 computer, later on acknowledged as a milestone in machine learning research. Newell and Simon were designing computer chess strategies, then turning to logic theorem proving: their hand simulation of



LOGIC THEORIST was completed in December 1955 (its first proof was printed by a JOHNNIAC computer in August 1956). Early computer simulations of perceptual tasks had been developed by Oliver Selfridge and Alfred Uttley. Computer simulation of neural nets, stemming from the seminal work by McCulloch and Pitts [11], were in progress, in particular by Farley and Clark [7], and by Rochester and some co-workers (including John Holland), regarding Donald Hebb's theory of learning and concept formation. In turn, both Minsky and McCarthy were dealing with several issues concerning machine intelligence.

The latter experiments are alluded to or mentioned in the Dartmouth preparatory document of August 1955. But another important event took place around that time: the Symposium on "The design of machines to simulate the behavior of the human brain", sponsored by the PGEC at the IRE National Convention held in March 21-24, 1955. The panel members were McCulloch, Anthony Oettinger, at the time at Harvard, Rochester, and Otto Schmitt, a biologist and an eclectic figure of science. John Mauchly, Marvin Minsky, Walter Pitts, Morris Rubinfeld were among the invited discussants. The transcripts of this less-known Symposium are enlightening (see [12]). They are a unique inventory of the main issues involved in the building of intelligent machines, of methodological approaches, ambitions and difficulties that would move to the forefront during the following decade, and in some cases even in more recent times.

One of the main issues dealt with at the Symposium was the possibility of using computers for different aims, and what might be the "the neurophysiologists' contribution" to the building of machines reproducing brain functions. In his talk "Contrasts and similarities" (see [12]: 242-242),

Oettinger distinguished two approaches in simulating human brain functions by computers, which, "although related, are far from being identical". The aim of the first, more engineering-based approach is the building of efficient machines per se, as aids in human intellectual tasks; the aim of the second, a more theoretically oriented approach, is the understanding of the human brain and behaviour. Here is probably the first clearly formulated statement of a distinction between two approaches in machine intelligence which was to become canonical in the AI community.

In the former, the more engineering-based case, the aim of simulation is to build computers that effectively duplicate or amplify human mental abilities. One might ask to what degree knowledge of the brain could be useful to the machine designer in this case. Oettinger's claim was that this issue is a controversial one. The designer might try to solve many computing and control problems using abilities in which the computer excels, e.g., speed and accuracy of computation, eventually trying to join these abilities with those in which human brain excels, e.g., degree of freedom, adaptability to new situations,

and so forth. But in any case, simulation deals with brain functions, not with brain structure.

Oettinger pointed out that most successful simulations of living functions had usually been achieved not by "following the example of nature", but by using structures and means not used by living organisms, thus attaining also superior performances of living functions: "for example, while the flight of birds undoubtedly stimulated man's urge to fly, human flight was achieved by significantly different means" (this is an example, by the way, which would become popular in the AI community afterwards). As for digital computers, on the one hand, their structural features are different from those of the human brain (Oettinger mentioned here John von Neumann's estimates regarding the reliability of the components of brain and computer), on the other hand, computers successfully perform arithmetic operations using processes different from those of humans, and it can be expected that "many machines of the future will continue to have only a functional resemblance to living organisms".

In the second case, the more theoretical one, the aim of simulation is quite different in Oettinger's view: computers are tools for testing hypotheses regarding brain functions, i.e. they can be used as neurological and psychological models. For Oettinger, two distinct cases are possible here. First, one has a theory of brain functions stated in mathematical form, such as Bush and Mosteller's theory of conditioned learning [2]. In this case, the computer can be used as in ordinary engineering applications, to solve differential equations, to obtain numerical values of functions, and so forth. Second, one has a theory stated so to speak in verbal form, as Hebb's theory of learning and concept formation. Hebb [8] introduced the notion of "cell assemblies", or nets of neurones strongly connected through excitatory synapses. As a result of repeated co-activation of constituent neurones, cell assemblies develop, as stated by Hebb's well known postulate. In this case, Oettinger concluded, "the digital computer may be programmed to simulate the neurone network with its environment", with the aim of testing Hebb's theory, as shown by the simulation program illustrated by Rochester at the Symposium, which I mentioned above among the early attempts in simulating neural nets.

Rochester presented a set of simulation experiments on an IBM 701 in his talk "Simulation of brain action on computers" (see [12]: 242-244). To put it briefly, a first simulation of cell assembly theory seemed to show that Hebb's postulate was not sufficient: co-activated neurones did not spontaneously develop cell assemblies. Further simulation experiments were carried out, based on a modification of Hebb's theory proposed by "one of Hebb's students", as Rochester said at the Symposium without mentioning him (it was, in fact, Peter Milner). Then a network of 63 simulated neurones, each connected to about eight others, was considered, and simulation tests



of a revised version of Hebb's theory were in progress at the time with better results. This exemplifies a general computer modelling methodology, here explicitly stated for the first time in the framework of the nascent AI, which would become pervasive in brain and behavioural sciences up to our time. It goes from formulating the model as a computer simulation of a theory of brain functions, to determining the implications of the model, to testing them, and finally to using data to prove, disprove, or modify the model, or the theory itself. (This machine simulation methodology has its own ancestors: see [4]).

Rochester's simulation methodology was positively evaluated by AI pioneers concerned with the aim of realistic simulation of human behaviour by computers, such as Newell and Simon. They believed that a brain or behavioural theory stated as a computer simulation model was, in general, the best alternative both to verbal (qualitative) descriptions of the theory, such as that originally given by Hebb, and to mathematical (quantitative) statements, such as that by Bush and Mosteller, mentioned also by Oettinger at the Symposium (see [16]: 396-397). According to Newell and Simon, to formulate a theory of human behaviour in terms of a computer program is to state an "Information Processing Theory" (for further details, see [5]).

### 3. Emboding flexibility in computers

At the 1955 Symposium, questions raised by Otto Schmitt in his talk "The brain as a different computer" (see [12]: 244-246) were much debated. As a biologist, he stated contrasts between the ordinary digital computer and the biological brain from a point of view different from Oettinger's. For Schmitt, computers should imitate the flexibility of reasoning usually shown by humans in order to be good simulators of brain function. Thus, computers would have to use a kind of loose or "grey logic", as he put it, not the rigid, bivalent, or "black-and-white logic" that presently characterises them. This would allow computers to grasp ill-defined and abstract concepts, as well as to exploit the incomplete, conflicting or partially inappropriate knowledge commonly available to humans in real life, e.g., in problem solving or decision making situations.

These rather vague statements took on a slightly more specific form in the discussion that followed Schmitt's talk. The issue could be so stated: how can common-sense knowledge be embodied into computer programs, as regards both complex and real-time human decision making? As to complex decision making, Schmitt judged, in replying to Oettinger, that programmers should seriously consider how to embody in programs those flexibility-based features of the brain, not only when the aim is "a realistic simulation of brain behavior" (as Oettinger put it in the discussion) but also when the aim is building efficient machines, as aids in human intellectual

tasks, i.e. "as tool[s] to do something for [us]" (again Oettinger). Even in the latter case, Schmitt concluded, "it is necessary to abandon the idea of perfectly correct, uniformly logical solutions in any machine which is to arrive at generally appropriate quick solutions to complex problems when provided only with sketchy, conflicting, and partially inappropriate information and instructions" (see [12]: 247). This is also true as to situations regarding quick, real-time decisions, as in Schmitt's example of a driver who might have to decide on exceeding established speed limits, given a particular road situation—this and analogous examples are presently proposed as instances of situated actions in AI (for a discussion, see [5]: chap 7). This decision is easy to make for a human, but it would be most difficult for a rigidly programmed computer. Thus the programmer should give the machine "a great deal of tradition and factual information, and some personal opinion", and an ability to revise its conclusions, a move not allowed in classic-logic based reasoning:

In this occasion, Oettinger was optimistic: "With computers it seems to me that we are able in principle, by the use of appropriate programming or designing of structure, to build in one swoop the whole background of explicit existing knowledge" (see [12]: 249). This seems a prelude to the future debate on how to embody abstract concepts in computer programs, and on the very possibility of grasping the background of explicit knowledge by them: an issue regarding what will be called the knowledge representation problem in AI. How to get a computer with common sense has been at the core of McCarthy's and Minsky's research, albeit from different points of view, since the very beginning of AI.

Some of the contrasts were stated by Schmitt quite vaguely and even improperly. Consider, for example, the contrast between a "systematic" (i.e. computer programmed) and "non systematic" (i.e. sketchily informed) search for complex-problem solution. Non-systematicity seems to include some kind of random elements of a not clearly specified nature—an issue touched on in the Dartmouth document [9], which, as with the transcript of the Symposium, includes a brief discussion of some search procedures including randomness, such as the Monte Carlo method.

One should notice that computer programs began to be capable of such non systematic search procedures. Schmitt's vaguely defined non systematic machine, able to get answers to problems "with feedback checks of results", is precisely the machine that Newell and Simon, with Clifford Shaw, were experimenting as the LOGIC THEORIST. This machine was endowed with a particular problem solving procedure (actually, a heuristic one) capable of "obtaining a feedback of the results [of a choice] that can be used to guide the next step" towards the solution (see [15]: 121).

Schmitt seems to share here the idea that computers were machines following logically unalterable



procedures, and thus not capable of modifying their behaviour under differing circumstances. Contrary to Schmitt's conclusions, on the one hand, later self-organising-system research would try to grasp certain nervous-system computational features he pointed out, and on the other hand, the "sketchy, conflicting, and partially inappropriate information and instructions" he alluded to as characterising human problem solving would become the core of heuristic computer programming in the field of complex decision making. Let us see the latter point in further detail.

Another Symposium is a case in point here, the one sponsored by the PGEC at the March 1956 IRE National Convention, thus a few months before the Dartmouth Conference. It was one of the first meetings devoted to "The impact of computers on science and society" [1], and most speakers came not only from the academic world but primarily from government and industry. The impact of computers on science concerned engineering, physics, chemistry and biology, as well as human sciences. The impact on society concerned different computer applications in data processing, mainly in industry and government, e.g., in management, defence and welfare. It would seem that the above discussions on the ability of new machines to simulate human decision processes are here converging on an applied research field: how "to seek more effective techniques and devices to assist us in managing and arriving at best solutions to our complicated and varied problems" (see [1]: 143).

At the Symposium, speakers agreed on the current limits of computers as to this goal, but also on the fact that computer capabilities were either underestimated or not fully appreciated, so that the computer was "a new tool with great and still unrealised potential" (the computers referred to were above all the ILLIAC and the SEAC computers). A common claim was that, up to the time, computers had been firstly considered as large calculating machines, useful in business applications (i.e. in "computations concerning money"), but less in those areas, regarding government and industry, in which complex data processing and optimisation procedures are involved. The tasks at the centre of the various talks were mainly information classification and retrieval, optimisation in complex decision making and planning. At the time, the prevalent techniques in assisting humans in such tasks were borrowed from Operation Research (OR, in the sequel). Computers played an important role in this field starting from World War II at least, and OR was explicitly mentioned at the Symposium, as well as its difficulty in dealing with data processing and complex activities involving information processing and planning.

In the Symposium the interaction between OR and AI can be vividly seen at its germinal stage. It is not by chance that the new-born expression "artificial intelligence" is used here perhaps for the first time publicly before the Dartmouth Conference. It was used by John Mauchly—one of the builders of ENIAC along with

Prosper Eckert—in his talk at the Symposium, in dealing with the issue raised by David Sayre, at the time at IBM and one of the authors of FORTRAN with John Backus. The issue concerned decision making procedures in complex problem solving and planning (scheduling of production, control of traffic in airline systems, and so forth).

As is well known, the expression "artificial intelligence" was introduced by John McCarthy in the 1955 document proposing the Dartmouth Conference. Sayre's name was present on the list, attached to the document, of the people whom the organisers of the Conference believed might be potential participants, as interested in the AI research project. At the Symposium, Sayre touched on the issue of machine intelligence, speculating about a way to endow a machine with what he called "something that approaches intelligence". This kind of machine might have satisfied certain of the above criteria questioning the flexibility requirements of computers, and appears to be endowed with that "self-improvement" ability which characterises the "truly intelligent" machine alluded to in the Dartmouth document [9]. Sayre, however, explicitly related such an intelligent machine to decision making in OR, when he suggested that complex activities or tasks, such as the aforementioned problem solving and planning, required "a rather different technique of machine use than we have yet developed". Given that no "exact procedure" had been evolved for solving these problems, the issue at point was "how to cause a machine, which has been given a fairly exact procedure, itself to amplify and correct it, constantly producing better and better procedures" (see [1]: 149). These "fairly exact" or "inexact" procedures were underlined by Mauchly as a mark of machine intelligence: "It is certainly true that many of us are interested in what has been given the name 'artificial intelligence'. This is indeed a field in which a great deal is going to be done, and there will be much influence on the future applications if we are successful in some of the endeavors which [Sayre] described as coming under 'inexact' rules, procedures, and applications" (p. 155).

The point at issue here is the ability of computers to make decisions, simulating human problem solving procedures, in order to assist humans in complex information processing, planning and decision making. At the time, it was Simon who would have pointed out the limits of OR techniques in dealing with such complex situations, where information is sketchy, and procedures do not guarantee optimisation in decision making. Developments of new, AI-based programming techniques were promptly applied in the field of management and decision making, where economics and psychology seemed to converge. As Simon viewed it, "AI was born in the basement of the Graduate School of Industrial Administration at Carnegie Mellon University, and for the first five years after his birth, applications to business decision making (that is OR applications) alternated with



applications to cognitive psychology” (see [18]: 5). This is a personal view of the origins of AI, but it effectively points out the role of early AI in evolving new techniques in data processing.

Briefly, for Simon, the model of the decision maker is not the omniscient economic man, the *Homo oeconomicus* of classic economics, who maximises his choice as predicted by the game theory. Endowed as he is with an ideal rationality, economic man is assumed to be fully informed on the problem domain or environment, as complex as this may be. In fact, this model is an extreme idealisation, too far removed from the actual decision maker, who is commonly dealing with complex, usually ill-structured problem domains, about which he is poorly informed. Another, more realistic model was proposed by Simon, that of the “administrative man”. This deals with both computationally complex and real-life problems, and, endowed as he is with a kind of “bounded rationality”, as Simon put it, is usually unable to maximise his choice, so using “satisficing” decision procedures, finally called heuristics [3]. To put it a bit crudely: disciplines initially involved in using these two different models of decision maker were, on the one hand, OR, based on common linear programming and probability theory techniques, and, on the other hand, AI, based on the new-born heuristic programming. Briefly, in Simon’s view it was OR’s failure in dealing with more human-like problem solving procedures that was the major cause of its early divorce from early AI (for details, see [5]).

To conclude, both Schmitt’s “sketchy, conflicting, and partially inappropriate information and instructions” and Mauchly’s “‘inexact’ rules, procedures, and applications” seem to state requirements then met by early-AI heuristic, often human-like, rules or procedures. As regards the above mentioned areas, concerned with management and complex decision making and planning, the background in which those requirements were initially met is the theory of the administrative man, developed by Simon starting from the 1940s.

#### 4. Conclusion

On the thresholds of the Dartmouth Conference, and against the backdrop of the spread of early large digital computers, several issues were raised that would influence both future research areas and future controversies in AI. To sum up, I would mention the following:

- how to use non numerical, i.e. symbolic, programming in the simulation of human abilities by machines;
- how to state different uses of computers: on the one hand, in realistic simulations of organism behaviour and, on the other, in efficient engineering and management applications;

- how to state the theory-model relationship within non-numerical computer simulation, given empirical facts and theoretical hypotheses regarding the brain or behaviour;
- how to justify the role of neurophysiology, having identified different levels of investigations—behaviour processes and brain processes—both considered, however, as functional levels;
- how to embody knowledge in computers, and what kind of logic would be useful, firstly as regards real-life situations;
- how to relate decision making and OR with new-born AI techniques—apparently more capable of dealing with complex and ill-structured problem domains.

Heuristic programming has been the case in point here. LOGIC THEORIST has been considered the first heuristic program. Although it played an important role at the Dartmouth Conference, programs different in complexity, such as Samuel’s above all, included procedures that could be called heuristics. Advances in earlier heuristic programming, also seen as a promising approach in data management and complex decision making, are among the most relevant causes that made program simulation of human behaviour prevail over distributed, self-organising and neural net approaches. These began to be rapidly and diffusely seen as a more brain-like style of computation, in particular when AI—as a new science of the mind—was suggested to be a level of behaviour explanation autonomous from the nervous system level (or levels).

Two years after Dartmouth, at the 1958 Teddington Symposium, the opposition between “imitators of the mind” and “imitators of the brain”, as Pitts put it, was definitively stated by Minsky, in his review of earlier advances in heuristic programming (see [4]: 187-189). Minsky [13] opposed hierarchic systems, “dealing with rather clear-cut syntactic processes involving the manipulation of symbolic expressions” to “‘network’ machines”, endowed with fairly simple self-organisational capabilities. He claimed his disaffection from the latter, if “really sophisticated behavior” is to be simulated. Moreover, it would not have been surprising if, once presently unknown nervous system mechanisms of intelligent activities had been identified, “the remaining heuristic theory would not be very different from the kind concerned with the formal or linguistic models”. At the moment, it might be thus worthwhile, Minsky concluded, to devote major efforts to heuristic programming, or what “some of us call ‘artificial intelligence’”.

#### REFERENCES

- [1] A.V. Astin, L. Cohen, J.W. Forrester, A.W. Jacobson, J. Mauchly, R. E. Meagher, and D. Sayre. The Impact of



- Computers on Science and Society. *IRE Transactions on Electronic Computers*, EC-5:142-158, 1956.
- [2] R. Bush and F. Mosteller. *Stochastic Models for Learning*, Wiley, New York, 1955.
- [3] R. Cordeschi. The Role of Heuristics in Automated Theorem Proving. *J.A. Robinson's Resolution Principle. Mathware and Soft Computing*, 3:281-293, 1996.
- [4] R. Cordeschi. *The Discovery of the Artificial: Behavior, Mind and Machines Before and Beyond Cybernetics*, Kluwer Academic Publishers, Dordrecht, 2002.
- [5] R. Cordeschi. Steps Towards the Synthetic Method. *Symbolic Information Processing and Self-organizing Systems in Early Artificial Intelligence Modeling*. In M. Wheeler, P. Husbands and O. Holland (Eds.), *The Mechanisation of Mind in History*, MIT Press, Cambridge, MA, 2007 (forthcoming).
- [6] D. Crevier. *AI. The Tumultuous History of the Search for Artificial Intelligence*, Basic Books, New York, 1993.
- [7] B. G. Farley and W. A. Clark. Simulation of Self-Organizing Systems by Digital Computer. *IRE Transactions on Information Theory*, 4:76-84, 1954.
- [8] D.O. Hebb. *The organization of behavior*. Wiley and Chapman, New York and London, 1949.
- [9] J. McCarthy, M. L. Minsky, N. Rochester, and C. E. Shannon. *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*, 1955. URL: <http://www.formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>.
- [10] P. McCorduck. *Machines Who Think*, Freeman, San Francisco, 1979.
- [11] W.S McCulloch and W. Pitts. A Logical Calculus of the Ideas Immanent in Nervous Activity. *Bulletin of Mathematical Biophysics*, 5:115-137, 1943.
- [12] W.S. McCulloch, A. G. Oettinger, N. Rochester, and O. Schmitt. *The Design of Machines to Simulate the Behavior of the Human Brain*. *IRE Transactions on Electronic Computers*, EC-5:240-255, 1956.
- [13] M.L. Minsky. *Some Methods of Heuristic Programming and Artificial Intelligence*. *Proceedings of the Teddington Symposium on Mechanisation of Thought Processes*, vol. 1, pp. 3-27, H.M. Stationery Office, London, 1959.
- [14] P. Mirowski. *Machine dreams. Economics becomes a cyborg science*, Cambridge University Press, Cambridge, 2002.
- [15] A. Newell, C. Shaw, and H. A. Simon. *Empirical Explorations with the Logic Theory Machine: A Case Study in Heuristics*. *Proceeding of the Western Joint Computer Conference*, pp. 218-239, 1957. Reprinted in E.A. Feigenbaum and J. Feldman (Eds.), McGraw-Hill, New York, pp. 109-133, 1963.
- [16] A. Newell and H.A. Simon. *Computers in Psychology*. In R.D. Luce, R.R. Bush, and E. Galanter (Eds.), *Handbook of Mathematical Psychology*, vol. 1, pp. 361-428, Wiley, New York, 1963.
- [17] C.E. Shannon. *Computers and Automata*. *Proceedings of the IRE*, pp. 1234-1241, 1953.
- [18] H.A. Simon. *The Future of Information Systems*. *Annals of Operation Research*, 71:3-14, 1997.