NATIONAL PHYSICAL LABORATORY

SYMPOSIUM No. 10

Mechanisation of Thought Processes

VOLUME I

NATIONAL PHYSICAL LABORATORY

SYMPOSIUM No. 10

Mechanisation of Thought Processes

Proceedings of a Symposium held at the National Physical Laboratory on 24th, 25th, 26th and 27th November 1958

VOLUME I

LONDON: HER MAJESTY'S STATIONERY OFFICE
1959

P R E F A C E

This Symposium was the tenth in the present series of N.P.L. Symposia. Two are normally held each year - one on a subject of general industrial interest and the other on a theme of more academic research.

This Symposium was held to bring together scientists studying artificial thinking, character and pattern recognition, learning, mechanical language translation, biology, automatic programming, industrial planning and clerical mechanization. It was felt that a common theme in all these fields was "The Mechanization of Thought Processes" and that an interchange of ideas between these specialists would be very valuable.

It is unfortunate that meeting accommodation in the Laboratory is at present very restricted, and a very large number of people had to be turned away. Nearly 200 delegates attended and of these about one third came from overseas.

A total of 32 papers was presented, and on the fourth day there were parallel sessions - one covering implications for Biology and the other for Industry. With one exception these papers are reproduced here substantially as presented i.e. with only very minor revisions. Several authors have added appendices to their papers since the Symposium. The discussion was recorded and all contributors and authors were asked to edit their contributions. The discussion is reproduced in full.

There were also a number of Lecture-Demonstrations, and a list of these is included. Most of the demonstrations are described in the papers or appendices, and the others are described in short papers.

THE MECHANIZATION OF THOUGHT PROCESSES

OPENING ADDRESS

by

DR. G. B. B. M. SUTHERLAND

THE National Physical Laboratory occupies a position midway between a university and an industrial laboratory. One of the ways of carrying out our function is to organise symposia in fields in which we are actively engaged and in which universities, industry and other Government laboratories are interested. That this symposium is typical is shown by the fact that roughly one third of the delegates are from universities, one third from industry and one third from government institutions. This was not an administrative decision - it just came out that way.

For the National Physical Laboratory to organise a symposium on "The Mechanization of Thought Processes" is very appropriate, because the future program of the Control Mechanisms and Electronics Division is closely concerned with certain aspects of this problem. Our general aim is to develop equipment which will carry out many of the tedious but essential mental tasks which at present are performed by large numbers of human beings, e.g., pattern recognition, retrieval of information, language translation, and the control of complex operations by trial and error learning. How important is it to solve these problems? We are sometimes criticised for spending time on making, or trying to make, machines, the performance of which compares very unfavourably with that of a human being on the same task, e.g., language translation. There are, of course, three answers to this criticism. The first is that the performance of a child bears little relation to that of the adult into which it develops. The second is that it may not be necessary to do the task as well as the human being, provided the machine does it much more rapidly or at much less cost. The third is that in trying to solve problems of this kind one inevitably learns a great deal about the essential principles which are inherent in the biological solution of the problem, but which would otherwise have lain undiscovered. It is becoming clear that the language of the control engineer has a place in the thinking of biologists and, more generally, that the growing interaction between biology, physics and electronics is likely to prove extremely fruitful to science in the quite near future.

(Dr. G. B. B. M. Sutherland, F.R.S., is Director,
National Physical Laboratory, Teddington)

The thinking process brings to mind Descartes' famous remark, "Cogito ergo sum" - I think, therefore I exist. We have as yet no glimmering of how a brain thinks about itself. Even after that stage is reached, we still have the problem of simulating a process carried out at the molecular level, by bits of electronic equipment. Our objectives this week are much more limited, but the problems are no less fascinating, and so let me invert Descartes' proposition and say that, since we exist, we can think, and let us now think to some purpose.

(94009)

CONTENTS

PAGE

SESSION I

GENERAL PRINCIPLES

Chairman: DR. A. M. UTTLEY, NPL

SESSION 1

PAPER 1

SOME METHODS OF ARTIFICIAL INTELLIGENCE AND HEURISTIC PROGRAMMING

by

Dr. MARVIN L. MINSKY

BIOGRAPHICAL NOTE

Marvin Lee Minsky was born in New York on 9th August, 1927. He received his B.A from Harvard in 1950 and Ph.D in Mathematics from Princeton in 1954. For the next three years he was a member of the Harvard University Society of Fellows, and in 1957-58 was staff member of the M.I.T. Lincoln Laboratories. At present he is Assistant Professor of Mathematics at M.I.T. where he is giving a course in Automata and Artificial Intelligence and is also staff member of the Research Laboratory of Electronics.

SOME METHODS OF ARTIFICIAL INTELLIGENCE AND HEURISTIC PROGRAMMING

by

Dr. MARVIN L. MINSKY*

SUMMARY

THIS paper is an attempt to discuss and partially organize a number of ideas concerning the design or programming of machines to work on problems for which the designer does not have, in advance, practical methods of solution. Particular attention is given to processes involving pattern recognition, learning, planning ahead, and the use of analogies or "models". Also considered is the question of designing "administrative" procedures to manage the use of these other devices. The paper begins with a discussion of what is meant by "Intelligence" and concludes with a section concerned with some techniques through which a machine might further improve itself by adding to its collection of problem-solving methods.

1. INTELLIGENCE

I feel that it would not be useful to lay down any absolute definition of "intelligence" or of "intelligent behaviour". For our goals in trying to design "thinking machines" are constantly changing in relation to our ever-increasing resources in this area.

Certainly there are many kinds of performances which if exhibited by a man we would all agree, today, require or manifest intelligence. But would we agree tomorrow? For some purposes we might agree with Turing *(ref.24)*, to regard the same performances in a machine as intelligent. In so doing we would be tying the definition of intelligence to some particular concept of human behaviour.

While such a convention might be useful in some kinds of discourse, its use in serious analysis is precluded by two serious faults. First, it

* The work leading to this paper was supported, in too many ways to cite individually, by the joint services of the U.S.A.

directly evades any concise specification of the kinds of behaviour we are looking for. Second, we can often find simple machines which in certain situations do exhibit performances which would be called intelligent if done by a man.

We are, understandably, very reluctant to confer this dignity on an evidently simple machine. Hence the conflict one would suffer in using this definition would threaten any descriptive value it might otherwise have.

In what situations are we less reluctant to attribute intelligence to machines? Occasionally, a machine will seem to be more resourceful and effective than one might expect from casual inspection of its structure. We may be surprised and impressed and we tend to remain so until through analysis or "explanation" the sense of wonder is removed. Whenever a system behaves as though it had more resources than were evident at "first glance" we react in this way, and this reaction is closely related to that involved when we judge a performance to be an exhibition of intelligence. But clearly this reaction depends also on the resources of the individual who is making the observation. The behaviour of any machine (as we use the term) is always explicable in terms of its past states, external contingencies, and the causal or probabilistic relations between them. Hence the significance of the observer's surprise in this; it can be inferred that the observer is not so good a mathematician that his first glance constitutes an adequate analysis of the situation. In the same way, our judgements of intelligence on the part of other humans are often related to our own analytic inadequacies, and these judgements do shift with changes in understanding.

We frequently find that a skill which seemed highly intelligent in others becomes much less impressive when we have learned the trick of doing it for ourselves. Indeed, we refer to many very complicated procedures as matters of "skill" rather than of intelligence apparently just because there happens to be a known method of instruction through which the ability can usually be acquired.

In attempting to design intelligent machines we are, in effect, concerned with the problems of "creativity". Many people are hostile to such an investigation, maintaining that creativity (or intelligence) is some kind of "gift" which simply cannot be understood or mechanized. This view can be maintained only through a constant shifting of definition. As soon as any process or performance has been mechanized or otherwise explained, it must be removed, with qualifications and apologies, from the list of creative performances. This part is perfectly reasonable; once a process has been mechanized one no longer needs terms like "creative" for its description, and we, too, remove it from the list of things to be accomplished. The weakness of the advocate of inexplicable creativity lies in the unsupported conviction that after *all* machines have been examined some items will still remain on the list.

Let us put it clearly then, that in exploring what we call "the artificial intelligence problem" we are not looking for any kind of closed solution to any question like "what is intelligence and how can it be mechanized?". The judgement of intelligence is more a reflection on what we understand than on what we, or machines, can do. Instead, we are searching for new and better ways of achieving performances that command, at the moment, our respect. We are prepared for the experience of understanding and the consequent reshaping of our goals.

2. PROBLEM-SOLVING

How do humans solve problems? To begin with we have to replace our intuitive requirements by reasonably well-defined technical questions. This may require the largest part of the intellectual effort, but we cannot dwell on the subject. A minimal requirement is that one have a method of discerning a satisfactory solution should one appear. If we cannot do this then the problem must be replaced by one which is well-defined in that sense, and we must hope that solution of the substitute problem will turn out to be useful.

In the best case we come equipped with an efficient *algorithm:* a systematic procedure which, given the problem as input, is guaranteed to produce a solution as output; efficient in that the solution will arrive within reasonable bounds on time and effort. But for new and interesting problems we don't usually have algorithms, or at least not efficient ones. At the other extreme we may know nothing about how to get a solution (except for the ability to recognize one). In such a case we have no alternative save to launch into an exhaustive search through the ensemble of potential solutions, e.g., the set of all proper expressions in our language. Random search is no better in general than systematic exhaustion, and may introduce the possibility of failure. It is tempting but irrational to look for a panacea in chaos. But in any case it is well known that for interesting problems exhaustive search is usually out of the question, even with the aid of the most powerful conceivable machines.

Normally, we are not motivated to attempt such problems. "Interesting" problems always have roots in areas which are at least partially understood. We usually have a good deal of partial information about how to get a solution. But this information may occur in fragmentary form: we may have some information about the "form" of a solution, recollections of similar problems solved in the past, general suggestions, hints, and the like.

We need to find ways of writing programs which will be able to use these fragments, or general advice, to reduce the amount of search to reasonable proportions.