E. Feigenbaum

Discussion of the paper

"Artificial Intelligence: Themes in the second decade"

by E.A. Feigenbaum,

delivered at the IFIP Congress 1968,

Edinburgh, Scotland, Aug. 8, 1968.

by

Saul Amarel

RCA Laboratories

Princeton, N.J. USA

I am in strong agreement with Feigenbaum's assessment of the present state of the art in A.I. and also with his views about the trends of work in this area and the nature of research problems that are now central and deserve more of our attention. I would like to add a few comments that are mainly intended to emphasize some of the points made by Feigenbaum.

First on the relationship between the subject matter of A.I. and work in "conventional" computer programming. It is important to realize that there is no sharp dividing line between the procedures of A.I. and the procedures of today's conventional software - both systems software and applications software. Systems software is concerned with problems of language analysis and interpretation (macro- assemblers, compilers) and with a variety of problems of control and optimization of resources that are becoming increasingly complex with the advent of time sharing. The procedures that are being developed and used for these problems strongly resemble in their overall logical structure - and also in their technical detail - to heuristic procedures for theorem proving and optimization of the type studied in A.I. research. Also, there are many application packages today in engineering and management that are constructed on the basis of a combination of systematic and heuristic methods (e.g. wiring and location of components in integrated electronic modules, stock cutting, route scheduling). Today's outstanding example of the application of A.I. ideas to an important "real life" problem is Feigenbaum and Lederberg's computer-based system for spectrometry.

It is becoming increasingly clear that the advanced procedures of A.I. and the procedures that direct today's "useful" work of computers lie on some sort of continuum. The key variables in this continuum are the amount of systematic knowledge available about a problem class, and the degree to which this knowledge can be efficiently exploited for the solution of specific problems in the class, the latter depends on the <u>form</u> of the available knowledge. At the one end of the continuum where amount of formal knowledge and its grade of utilization are high we have most of today's "conventional" programs. At the other end, where the amount of systematic knowledge and its grade of utilization are low, we have the general, flexible, only partly validated procedures of A.I., where a set of relatively weak problem-specific principles are bombined with several powerful heuristic methods for organizing search processes.

One of the important goals in the development of problem solving procedures in A.I. is to enable a user to specify directly his <u>problem</u> to a computer - in its "initial" high-level functional form - without having to specify to the computer an explicit procedure for solving it. This possibility would be a major step in the road towards programming automation. It will bring the vast information processing power of computers much closer to the man-with-the-problem, and it will permit the application of computers to a much larger domain of intellectual tasks. However, even in this case the man is left with the responsibility of formulating his problem to the computer in a manner which promises a reasonably efficient solution-finding process. Here we are confronted with the problem of problem representation, which, I fully agree with Feigenbaum, is today's central problem of A.I. In representing a problem to a machine the man provides at present all the knowledge about the problem that the machine can work with, and also he provides it in a specific form which reflects his specific point of view - a point of view which may or may not be fruitful for the solution-searching process that he is forcing upon the machine. The question arises naturally whether it is possible to endow machines with capabilities to shift problem representations in an "appropriate" direction. Such a capability will indeed provide us with problem solving machines that combine great generality and power. I have been concerned with this problem in the last few years, and I feel at present that in order to realize beneficial shifts in problem representation we need to know more about the following two general questions.

(1) How to choose the basic concepts for a language in which problem situations, rules for transitions between situations, and general knowledge about the problem can be expressed. This is of particular importance in "real life" problems, where the problem is not formulated at the outset within a formal system but it is given verbally or it includes information obtained from physical sensors. I think that this is the fundamental problem in the work on "robotics", i.e., how to formulate descriptions of a physical environment - among the multitude of possible descriptions - that are most appropriate for the tasks on hand. The question of choice of descriptions for problem-relevant knowledge is also fundamental for the design of question-answering

systems with complex data bases. Choosing descriptions for data bases in information systems is certainly an important part of the problem of problem representations; in my opinion, this is a problem that deserves much attention, and its study is likely to produce many fruitful results in the art of computing.

(2) How to proceed in the discovery of useful properties of a problem space that can be used to transform it into a space where search for solution is less difficult, and how to use this knowledge in the formulation of a better problem solving procedure. This involves the detection of irrelevancies and redundancies, the recognition of regularities (such as symmetries) in the space, and the ability to form more powerful rules of action (say formation of macromoves from moves) that incorporate the newly discovered knowledge. It is conceivable that the formation of more powerful rules of action on the basis of new problem-specific knowledge is mechanizable with ideas and techniques available at present. To obtain non-trivial advances in this area we must know more about problems of formation type.

The question (1) and the knowledge-creation part of the question (2) are outside the realm of machines at present. However, I think that it is important to start exploring them with a view to possible mechanizations, via appropriately chosen case studies.

As a last comment I would like to indicate that most of the progress (in technique and theoretical understanding) in heuristic problem solving to date has centred on problems of derivation type, where the objective is roughly to construct a path between given boundaries. (e.g., theorem proving problems). Problems of formation type have received so far less attention. These problems are more difficult than derivation problems, and they involve reasoning from possible solutions to the problem conditions. Many "real life" problems, notably design problems and diagnostic problems, are of this type. Feigenbaum's spectrometry problem is largely a formation problem. Many problems of shift in problem representation are of formation type. I think that as we move more and more into useful applications of A.I. to complex problems, and as we attempt to attain more problem solving generality via computer handling of problem representations, we shall have to do much more work on formation problems.