Intelligent Design versus Solving Artificial Problems

Paper presented at the 1993 IEEE Region 10
International Conference on'Computers,
Communications, Control and Power Engineering' in
Beijing (China), October 19-21, 1993

Gerhard HANAPPI, Austrian Academy of Sciences, Vienna
Edeltraud EGGER, University of Technology, Vienna

INTELLIGENT DESIGN VERSUS SOLVING ARTIFICIAL PROBLEMS
ARTIFICIAL INTELLIGENCE RESEARCH AT THE CROSS-ROADS

Gerhard Hanappi and Edeltraud Egger
Academy of Sciences, Socioeconomics, Postgasse 7, A-1010 Vienna
University of Technology, CSCW, Argentinierstr. 8, A-1040 Vienna

ABSTRACT

The major hypothesis of this paper is that artificial intelligence research has reached a point, where two diverging future paths could be followed: On the one hand one could take the development of technical possibilities and try to find problems that fit to existing technical solutions. We call this route "solving artificial problems". On the other hand we propose an attractive field of future AI research, which concentrates on the design of empirically relevant solutions for man- machine interactions by splitting decisions in an intelligent way - a term to be defined explicitly in the paper - into those taken by machines and those taken by human groups. Using our notion of intelligent design we will give some examples to illuminate our case.

INTRODUCTION

Artificial intelligence research has reached a state of maturity which allows to look back at its roots from a more distant and critical point of view. It seems to be urgent to put forward such an evaluation since for many of the researchers involved the current situation is not satisfactory: The double trap of working either on the highly intelligent support of solutions for irrelevant problems or on the intellectually rather trivial support given to relevant problems has frustrated them. Had aspirations been too high or is there something basically wrong with AI?

In our view three trends contributed to the genesis of AI:

- A technical trend: Growing use of ever larger databases with more and more sophisticated logic incorporated in queries.
- A methodological trend: The approach of Newell and Simon to problem solving, in particular their idea of a general problem solver (GPS).
- A trend to name new research fields in a way that promises the allocation of research funds. A trendy label for the new field, namely 'artificial intelligence', was apt to attract financial support due to the enormous range of applications implicitly announced in these two words.

It has been the combination of these three trends which caused the tremendous success of the new research area. And it will be the same components which will draw the limits for further developments of AI research.

This paper argues that the borders set by technology, methodology and usefulness do not lead in the same direction, future efforts are not likely to follow all three trends at the same time - AI research at the cross-roads. To discover which roads are open and what type of research they imply we will start with a brief description of what we think is the AI mainstream research program.

SOLVING ARTIFICIAL PROBLEMS

By far the easiest trend to predict, though hard to achieve, is the technological development of ever growing knowledge bases with faster access, elegant user interfaces and sophisticated but handy logic incorporated. Expert systems, the commercial little brother of AI research, surely will be able to exploit these features. Clearly the ability to provide an enormous amount of quick answers to well-defined queries gave and will give expert systems the aura of 'intelligent' systems. It is mainly the quality of inference, of search if one prefers a more modest expression, which distinguishes high-brow AI from more commercially oriented expert systems. But inference machines too will surely benefit from 'more space and less time' as provided by new technology.

New hardware breeds new software, but it has to be doubted if the basic methodology incorporated really changes, or if improvements only concern the surface of applications. In our view the fundamental methodology of AI did not change too dramatically since the days when 'general problem-solving' was formulated as search in the early sixties. In a recent contribution[1] Newell and Simon restate and assert their view that the 'ability to solve problems is generally taken as a prime indicator that a system has intelligence' and that problem solving is done 'by generating potential solutions and testing them - that is by searching'. The link to technological possibilities is evident: If intelligence is indicated by problem-solving capacity and if the latter is defined as search, then these narrow definitions imply that better search algorithms in larger databases are equivalent to more intelligence. One has to take a closer look at Newell/Simon's lucid but reductionist notion of problem-solving to discover the methodological impasse they propose.

According to them problem-solving consists of two steps: generating a possible solution, a pattern, and testing it, comparing it to acceptable patterns. The generating part means running through a finite set of patterns, while the test simply provides an end condition. It seems to be a rather heroic assumption to state that this procedure lies at the heart of what we call 'intelligence'. To be more specific: We would insist that neither the range of possible solutions nor the set of acceptable patterns could be taken as given when talking about 'intelligence'. Quite the opposite should be assumed: 'Intelligence' has to be attributed to an entity, which can determine a solution space and goals, which is capable to specify problems (compare our work on decision support systems[2]). If the problem is specified and goals are known the pivotal part of 'intelligent' input has been accomplished. Of course there usually is a lot of non-trivial work to be done - it is not evident how to win a game of chess though rules and goal are given - but computer support given to these subordinated activities should not be mistaken as representing 'intelligence'.

Thus the claim to produce a general problem solver stems from a misspecification of the notion 'problem-solving'. Search algorithms can be general, problem specifications cannot - they are always particular. Of course the particular problem specification process can and will be supported by computer software. But 'intelligence' could only be a feature of the entity carrying out the whole process and not of one of the tools used. Clearly this critique not only applies to the 'symbolist school' of AI researchers but also to the 'connectionist school'. As one of the protagonist of connectionism, M.Minsky[3], put it: 'This broad

division (between symbolic and connectionist, G.H./E.E.) makes no sense to us' - at least from a methodological point of view.
Coming back to our main argument we therefore conclude that mainstream research in AI has been driven by technological improvements of databases and search methods, which have built on the same unchanged methodological paradigm sketched above. It is this underlying paradigm which now has lead to a certain frustration: Many of the early promises of AI, so useful to attract attention, would now have to be discarded. If the mainstream approach encounters insurmountable difficulties with its original tasks, then it has to look for new tasks: Artificial problems for artificial intelligence. The problem with 'real life'-problems is that almost none of them is well structured and ready for heuristic search of optima. Artificial problems either come from games or areas which resemble games, like psychology, law and the like. To enter the area of problem specification the fundamental methodological premise that intelligence can be reduced to search would have to be revised.
This, of course, is not likely to happen. To much financial and intellectual effort has been put on this project. The broadest and most probable road in the future is to carry on 'business as usual'. Solvable problems will be singled out artificially to prove the success of the prevailing paradigm. Though we think that this road is an impasse, we are sure that the scientific community in the area will not be easily to convince.

INTELLIGENT DESIGN

At the outset of our alternative to this first path for future AI research we have to clarify the relation between entity and tool. Entities (individuals, groups, firms, nations, society) have goals which they try to achieve. In the process of goal achievement they face certain difficulties which they try to specify as problems, that is as mental counterparts of the real difficulty. In doing that entities use different tools: pencil and paper, discussions, computers, etc., and they can perform more or less successful. To a certain degree their activities can be judged to be intelligent. Though it is hard to define a quantitative measure for intelligence, from a qualitative point of view it is clear that there exist several characteristics, Schank[4] mentions seven of them, which have to be considered as attributes of an entity if one wants to evaluate its intelligence. Proper use of language thus demands that 'intelligence' is a property which can be attributed to entities and not to tools[*] . A plan put forward by an intelligent entity often will be called intelligent too, but disregarding the entity under consideration and calling a tool intelligent turns things upside down: 'Users' start to appear as tools of intelligent machines. From a marketing perspective for expert systems this is straight forward; to praise a product with words usually restricted to human academics is quite clever. From a scientific point of view this confusion caused a lot of unnecessary misconceptions (e.g. Minsky[6], Morawec[7]).
The concept of a tool is quite interesting in another respect too: One entity can use other entities like tools. A firm can consider groups of workers like tools for its goals, a state can consider industries in that way too. This is a basic source of conflicts: Considering an entity as a tool means subordinating it to one's interests. Since subordinated entities do have their own interests - which can be contradictory to the more powerful entity, a permanent conflict potential is given. Conflicts appear as soon as one of the involved parties tries to

[*] To call algorithms intelligent, which behave in certain respects like intelligent entities restricts the concept of human intelligence to those respects: 'Many intelligent activities besides numerical calculation and information retrieval have been accomplished by programs.'[5]

realize its goals at the expense of the other one. We would argue that this situation is the standard case in the everyday-life of organizations.
To make things worse, not only diverging goals of interdependent entities lead to conflicts, these goals themselves are dynamically changing bundles of expected utilities. Entities permanently have to build (with relative utility) weighted averages over prospective outcomes of relevant future processes. New goals might be discovered, interdependencies between goals might be exploited, the goal finding process of negotiators could be intentionally disturbed, a variety of more or less costly actions could be taken. As a consequence the very process of goal formation itself is a candidate for computer support.
Having come to grips with one's goals is only the first step. The next step of course is to choose the appropriate actions to reach these goals. Again the whole host of uncertainties mentioned above re-appears. As a matter of fact this is the research subject of mathematical game theory, which by the way did not receive adequate attention in computer science. Even though there exists a wide range of simulation tools in computer science it has been neglected to go beyond the treatment of physical systems towards social systems. And social systems indeed are not a trivial application field. Note that such a type of computer support will not necessarily be the same as the one used for goal formation.
The results of punctually simulated possible goal realization processes then can be matched with the original aims. This third step is the one where the approach critizised in chapter 1 starts its inquiry: The search for optimality. Contrary to what has been the objective there, an overall solution, in this 'application-driven' approach the result is the most advantageous solution for a certain entity in a running conflict. Although the same algorithms might be used, in this second approach 'search' is just one component of computer support among others.
Since conflicts are conflicts between entities, they have to be carried out by entities. Negotiating and bargaining, as the most civilized way to do so, again can be given some support.

This seems to be the field proper of the rapidly developing research area of CSCW (Computer Supported Co-operative Work). It ranges from communication support to the use of simulation, proposing organizational changes as well as technical solutions.
Negotiating, of course, cannot guarantee that conflicts are solved. After each round of negotiations the groups involved will come back and reconsider their position - they will try to learn. Support of learning certainly is one of the most exiting areas of AI research. The tools which emerged - namely pattern recognition, the different methods machine learning, even econometrics - are waiting to be used for negotiation preparation[*] .
Since the resolution and re-appearance of conflicts is an ongoing process, the above mentioned steps have to be iterated. Intelligent design of computer support for social processes means that users are provided with a framework picturing power-relations and their potential dynamics (that is conflicts). As contradictions cannot be solved by a machine, intelligent design also includes assessment of switching points between user and machine: A system designer therefore has to start off with an overview of the organization identifying the conflict points and 'critical' decisions[*] which are the basis for face-to-face negotiation processes.

---

[*] As a recent source for the use of AI-techniques in economics see Moss/Rae[8].
[*] 'Critical decisions' are the ones which concern one's power position within the hierarchy of the organization[9].

CONCLUSION

The intelligence of an entity is mirrored in its design. A system designer trying to improve an existing design evidently will have to use many heterogeneous tools, this was what we wanted to make plausible in the last chapter. Some of these tools might come from mainstream AI and need adaptation, others exist just in nucleus, others do not exist at all. Application-driven design means to combine a set of newly developed or adapted tools according to the inner logic of the application. This includes a definition of the limits of the respective support as well as a decision on which of the conflicting entities within a larger entity* is be supported to which extent. In other words, the designer takes part in conflicts by structuring them and allocating support for negotiators - she/he becomes part of the larger entity, part of its intelligence.

Seen from this perspective the methodology of application-driven intelligent design differs radically from the one followed by mainstream AI research: Intelligence looses the character of being something which can be abstracted from entities and thus can be formulated as generally applicable algorithm. It is a naturally evolving property of entities, a process in which computer scientists play an increasingly important part. The challenge in the second approach is not to discover a general representation for semantics, but to integrate and to synthesize work done in different fields inside and outside computer science to arrive at designs accelerating the development of organizations.

---

* What we called 'inner logic' is the processing of conflicts between smaller entities, say workers and management, within a larger entity, say a firm.

REFERENCES

1. Newell A and Simon H, In: Boden MA, The Philosophy of Artificial Intelligence, Oxford University Press (1990), pp.105-146.
2. Egger E and Hanappi G, From Syntax to Semantics and Back: The Problem of Information Selection in GDSSs, Proc. of HCI International '93, Orlando, Florida (1993).
3. Minsky ML and Papert S, Perceptrons, MIT Press (1988), p. IX.
4. Schank R, Tell Me a Story, Macmillan (1990), pp. 219-237.
5. Barr A and Feigenbaum E, The Handbook of Artificial Intelligence, Vol 1, Addison-Wesley (1981).
6. Minsky M, Society of Mind, Simon&Schuster (1985).
7. Morawec, Mind Children, Harvard University Press (1988).
8. Moss S and Rae J, Artificial Intelligence and Economic Analysis, Edward Elgar (1992).
9. Egger E and Hanappi G, Competence Switching Managed by Intelligent Systems, Proc. of ISMIS '93, Trondheim, Norway (1993).