REAL AI

Matjaž Gams
Jožef Stefan Institute, Jamova 39, Ljubljana, Slovenia
Phone: +38 61 159 199, Fax: +38 61 158 058
E-mail: matjaz.gams@ijs.si

Keywords: artificial intelligence, viewpoint

Edited by: Branko Souček
Received: December 12, 1992 **Revised:** February 9, 1993 **Accepted:** March 1, 1993

In the first part, four viewpoints on AI are presented. It is proposed that a program exhibiting AI is one that can change as a result of interactions with the environment. While no program can be proclaimed as intelligent at the moment, intelligence may be just an emerging property of successful information machines or beings. In the second part, ideas, problems and misconceptions about AI are analysed through grouping into three categories: (1) facts - opinions that are supported by facts; (2) legends - opinions, based on facts but largely exaggerated; (3) myths - opinions not based on facts.

1 Introduction

Discussions about AI have attracted most of the researchers inside the field as well as many outside (Fox 1990, Hayes-Roth & Fikes 1991, Hayes et. al. 1992, Mettrey 1992, Schank 1991, Searle 1982, Wilkes 1992). In particular, there have been important shifts and modifications in world-wide opinion observed in recent relevant publications. These debates have motivated us to make an attempt to summarise them in a coherent way. [1]

First, let us analyse the notion of artificial intelligence.

2 Viewpoints on AI

Thinking about the definition of AI, one should ask 'Where's the AI?' (Schank 1991). There seem to be at least four prevailing answers to that question.

The **first view** sees AI as something magic that emerges out of a computationally effective computer after you put in it enough things. Indeed, it is still often acclaimed that neural networks mimic the behaviour of human brains. That is rather surprising since what they do at present is to mimic a numeric filter at best able to tune predefined parameters. Not surprisingly, computationally more diverse and structurarily more complex statistical methods typically achieve better classification accuracy (Henery & Taylor 1992).

This approach has already contributed to the first dark age at the beginning of AI and is still present in many subfields of AI. On the other hand, one should not underestimate their advantages such as flexibility and robustness.

The **second view** sees AI as a superb inference engine and therefore resembles knowledge engineering. What one has to do is to find an expert, encode his knowledge into lists of rules, add an inference engine with appropriate interface and there you are. This has led many people to think that AI means rule-based expert systems, and then they thought they understand them as well as AI. And since they have also learned the limitations of rule-based systems, they also think that is the limitation of AI, not just of one component of AI (Hayes-Roth & Fikes 1991). While connectionism as well as knowledge engineering and inference engines are important parts of AI research and applications, labelling it intelligent or as AI itself is misleading (Schank 1991).

The **third view** maintains that, if no machine ever did it before, it must be AI. For example, years ago research in computer chess was one of

[1] A similar but shorter paper (Gams 1992) has been presented as the first paper in the AI section at the ERK'92 conference - similar parts are reprinted with permission.

the central themata in AI. People thought that computer programs playing chess well would certainly have to be intelligent. Now, nearly everybody agrees that these programs do not have any deep intelligence at all. Luckily, there are several explanations of this contradiction. For one thing, this viewpoint seems to confuse getting a machine to do something intelligent with getting it to be a model of human intelligence. More important, if you define AI in that way - to be one of frontiers of computer science, once the area that you are looking at is understood, then it is no longer at the frontiers of computer science and therefore no longer AI, and so it is a no-win kind of situation (Schank 1991).

So, where is the intelligence in computer programs? As mentioned earlier, many difficult problems which had long been thought to require real intelligence, have been solved by rather unintelligent methods. Intensifying this argument, even superb intelligent behaviour does not guarantee that real intelligence and understanding have been achieved. For example, Searle (Searle 1982) has constructed a hypothetic Chinese room in which a group of workers performs intelligent translation between two natural languages (English-Chinese) and each of them performs only a subpart of the whole process on the basis of a predefined procedure. Although such a Chinese room could pass Turing's test, the room (and nobody inside it) does not understand the whole process and there is no real intelligence at all.

After a decade of quite intensive debate, there has not been any definitive answer to this paradox since it is actually a philosophical question: Is performance, i.e. a mechanicistic approach really sufficient? From a practical point of view, most things in our world work in the mechanicistic mode. Of course, there are paradoxes and unsolved questions (e.g. are there other universes or is there only ours?) but people have successfully lived with them. Not to mention that other approaches based on ideology or spiritism have not yielded similarly good results.

Therefore, it seems rather surprising that criticism is so strong in the AI area even if the same arguments are being repeated over and over again. A recent example of this kind could be the article "Artificial Intelligence as the Year 2000 Approaches" (Wilkes 1992). It provoked harsh replies (Hayes et. al. 1992) in which several errors and misconceptions were exposed. Nevertheless, in all this argumentation there are at least two points where the AI community still has to prove itself:

- if intelligence (in computers) were simple, fast and powerful computers would have facilitated it a long time ago, and
- many of the ideas in the AI field have produced much more optimism than real improvements.

Here we shall devote attention to the first argument, and the second argument will be analysed in the second part of this paper.

For example, Wilkes (Wilkes 1992) claims that intelligence may come from analogue circuitry since, obviously, it has not come from digital computers so far. Searle (Searle 1982) claims that digital machines can not be intelligent as biological beings since they are essentially different. Although Hayes (Hayes et. al. 1992) claims that no proof is given for such claims, the same is valid also for the reverse claim.

At this point we can only agree that real intelligence in machines has not been achieved yet. Furthermore, we still do not have any good ideas how to make a true intelligent machine. However, two arguments seem plausible:

- Real human intelligence is very complex. If it were simple enough for us to understand it, than we would be too simple to perceive that (as claimed by several authors).
- Intelligence may be just an emerging property of successful information machines or beings. There does not have to be any deeper motive or principle behind it. This approach is very close to the "artificial life" where computational models share many characteristics with biological computation (Brooks 1991).

Furthermore, in computing there are good foundations and clear concepts like Turing's machine or Church's thesis. There is also Turing's test in which real intelligence is achieved when human judges can not distinguish between the performance of a computer and human. Since computer programs are far away from achieving such a level, the contest area is often limited to a domain which still requires intelligence by human counterpart. However, it is important to notice that the communication between the judge and the contenders is an open one. Therefore, a computer program playing superb chess but unable to explain the motives of its moves certainly would not pass the test while even a novice player with normal explanation and reasoning capabilities certainly would. In recent years slightly modified tests, or competitions, are becoming annual events with rewards up to $100,000 (Epstein 1992).

This leads us to the **fourth view** on AI. True intelligence, exhibited by computer programs, would have to have many or even most properties of human intelligent behaviour regardless of how narrow the application area was. One of the main such properties is *learning* since intelligence first of all means getting better over time. In relation to Turing's test, a computer program unable to learn from its mistake would certainly be exposed. Today, hardly any AI programs learn from their mistakes, although - with very good reason, learning is the central area of AI at least in the last decade.

There is some additional reasoning about intelligence:

- Intelligence is in size. It is hard to expect a small program to display intelligence. Intelligence is neither simple nor easy understandable.

- Intelligence is in complexity and heterogeneity. This area is sometimes related to multistrategy learning (Michalski & Tecuci 1992), a multiple-knowledge approach (Gams et. al. 1991) and multi-agents (Minsky 1987).

- Intelligence is in the ability to perform well real-life tasks which require the use of knowledge. For example, Mathematica, a program for symbolic computing is regarded as approximately as intelligent as a numeric library. Contrary to recent interest in logic programming, it is quite probable that intelligence there will be at a similar level to Mathematica until real-life knowledge is incorporated into programs.

Furthermore, there are several aspects of intelligence each of which can be compared if not measured on a scale. For example, motional intelligence can be quite high in many animals. In another aspect, AI research can well be at the frontiers of computer science while AI applications fell into an application area years away from scientific achievements. AI applications do not have to be intelligent, they have to be related to AI research similar to other science/application relations.

In short, while agitated debates about AI raise interest and in both ways affect funds, what really matters is what works and which new discoveries are produced. *It is not that AI needs definitions; it is more that AI needs substance* (Schank 1991). Although general artificial intelligence has not yet been achieved, we know more and more about it. Some basic facts, legends and myths about AI will be represented in the following sections.

3 Facts

The first AI concept is search. Most difficult problems involve choosing between alternative solutions and evaluation processes in which the best solution is found. This basic search schema may not be immediately observed in diverse subareas of AI such as scheduling, games, learning or expert systems. Novice readers in AI might get distressed by different terminology and diverse techniques. However, even one of the oldest definitions of AI promotes it as a fight against combinatorial explosion.

While faster computers certainly help, simple search techniques can not ever deal with the exponential growth rate of the number of possibilities in a search tree. For example, in a single factory having 85 orders, 10 operations, and only one substitutional machine, one could create over 10^{880} alternative schedules (Fox 1990) while the number of all atomic particles in our universe is estimated at 10^{80}. Obviously, the key question is how to reduce search space.

The second AI concept is knowledge representation. It is not that the knowledge representation concept is second to search; it is one of the two. Knowledge is probably even more important than search in biological systems. In real life, response to a specific pattern is usually prestored - learned through experience. This resembles the fourth viewpoint on AI presented earlier. But from a practical point of view, computers as well as AI were more successful in search than in knowledge representation.

Although there are many techniques from semantic nets to frames, the most successful AI applications so far are expert systems. At times, it was thought that the expert-systems approach enables a uniform solution to knowledge representation problems. It has led to overenthusiasm and overselling the technological possibilities. Now we know that expert systems are appropriate only when problems are relatively small and stable or can be decomposed into such subproblems, meaning that experts agree with each other upon a proposed knowledge base.

The main problem, how to represent different kinds of knowledge, complex and heterogeneous knowledge, and combine them into one system has not been solved yet. As a consequence, successful learning from interactions with the environment has not been, and quite probably can not be achieved without it.

AI copes with the search combinatoric explosion by using knowledge. The use of knowledge enables successful pruning of a search tree. For example, in an expert system OPEX (Gams et. al. 1991) for generating appropriate machining operation sequences, designed in cooperation with researchers from the Faculty of Mechanical Engineering and Jožef Stefan Institute, there are three levels of rules:
(a) Rules for applying basic machining operation. Example:
operation drilling : if
  gdb:fc is-a-cylinder-in and
  gdb:dc included interval(3,40) and
  gdb:lc/max(gdb:dc) = 10 and
  gdb:nc subset interval(11,12)
then
  fc := is-a blank and dc := 0 and nc := undefined
(b) Rules that define various possibilities of linking basic machining operations within an individual feature. Example:
from boring to drilling if true end.
(c) Rules for combining operation sequences that define which operation sequences should be adopted for a combination of features. Example:
combination drilling and drilling if true end.

The task of OPEX is to design operation sequences for a machine and a specified part, and sort them according to predefined criteria. Naive combining of operations quickly leads to combinatorial explosion, but through smarter selection of possibilities, i.e. by utilising domain knowledge, the combinatorics is reduced to a feasible level.

As indicated by previous example, AI enhances search by reformulating problems, through the use of opportunism, heuristics, and by abstraction and differentiation of quantitative models. These are techniques behind the general principle of using knowledge to control search. In essence they perform similar improvements of search as hierarchical search or dynamic programming, however, the use of knowledge can greatly improve performance.

AI systems can increase productivity. Various reports estimate the number of AI systems regularly in use to around 3000 with some of them being in use for more than 10 years. The main problem with such estimates is where to put borders between AI and non-AI applications. For example, is Prolog interpreter an AI application or not? Clearly, there is no intelligence in it. On the other hand, Prolog as well as Lisp and many other products were designed as a by-product in AI research. In our opinion, they should be included as AI applications as well as neural networks. Actually, marketplace AI-software packages fall into at least four major categories: programming languages, programming environments, problem-solving shells (for a class of problems), and application shells (specialised for a given domain).

For example, in our rather small country of Slovenia, in two AI laboratories at Jožef Stefan Institute and the Faculty for Electrical Engineering and Computing, around 60 applications were successfully performed in recent years and 5 original programming systems with several thousands of lines are in regular use (Urbančič & Křižman 1991).

4 Legends

AI systems are easy to build. Indeed, under specific conditions, improvements in speed and productivity are enormous when using AI systems. For example, having stored a history of events, it is possible to design an expert system with the use of inductive learning tools in a couple of days. On the other hand, there are problems which take more or even much more time than by classical methods.

Specifications and prototyping largely enhance productivity. This is partially true. If the problem fits an application shell, knowledge gathered from experts can be put into a system quickly and then tested. Rapid prototyping elicits the requirements and specifications of software for ill-defined problems; in recent years it has been included in software development approaches as another example of AI product finishing in classical computer science and applications. But the limitations of the methodology and conditions for successfulness have also become known.

AI systems are easily verified and maintained. Since AI systems rely on knowledge instead of formulae, e.g. expert knowledge in expert systems, it is often propagated that these systems are highly understandable and, therefore, easy to be verified and maintained. For example, expert systems provide explanation possibilities as a sort of rule tracker instead of 'trace' in conventional programming languages. Practical experience has shown that while it is an important improvement over classical methods, verification and maintenance remain time consuming phases.

5 Myths

Artificial intelligence approach does not need conventional program-engineering and management techniques. This incorrect belief is still quite common due in part to academic ignorance of the requirements for building production-level systems.

Systems working on simple examples can easily be upgraded to full-scale real-life systems. Performing speech understanding for a small vocabulary of, say 50 words, differs greatly from the same task but with thousands of words. Similarly, many problems are difficult only because of their size. The myth of simple scaling is still very alive mainly due to an academic approach where it is most important that idea is fresh and attractive (working on a simple, carefully designed problem). Literature reviews in AI show that about half of all publications belong to this category and only half of the systems actually work on non-toy problems. In the worst scenario, some subareas of AI have for years attracted interest and funds without actually producing a program working on a realistic problem. There seem to be certain similarities to fashion movements in which a new direction promoted by famous people attracts global interest. After a critical mass is obtained, the movement can sustain for several years without any realistic verification. The problem is similar in several other sciences. The "publish or perish" science tends to produce famous writers instead of famous scientists, researchers, engineers or inventors. However slowly, in AI it is changing in favour of more strict verification of results. For example, there are several projects which for years have evaluated different methods (Henery & Taylor 1992). Even at our laboratories we have been testing all available inductive learning systems for 5 years and making the results public.

Small systems can exhibit full-scale human intelligence. In serious AI circles it is known that it is not possible to simulate full-scale human intelligence without huge and complex systems and that searching for a genuine simple algorithm is similar to searching for perpetuum mobile.

If we have an expert, then we can create an expert system. Obviously, a lot more is needed; first of all a feasibility study.

AI does not need business motivation to produce valuable results. Several studies have shown that those initiated by management have a better chance of returning profit.

AI tools can enable novices to develop expert systems. Inexperience and lack of skill can not be compensated in any field.

Expert systems consist of expert systems. Typically, in expert systems there is much more than that, including lots of classical programming.

Expert systems perform as specialised, stand-alone programs. Actually, they access databases, conventional programming languages, operating systems etc.

All AI tools are the same. There are different categories.

All expert systems are rule-based. Many, but there is much more.

Expert systems do not make mistakes. In real life there is no such thing.

AI replaces conventional approaches. Rather, they can both be useful depending on conditions, and are often combined together.

AI knowledge engineering is all we need to know about AI. The more you know the better. Again, AI consists of many diverse subareas.

AI tools are good only for AI applications. AI software supports qualitative and quantitative reasoning equally well.

In simple expert systems an exhaustive search can provide solutions. Yes, for toy problems.

Tools equally support both forward and backward chaining. At the expense of the other.

The more general the tool the better. Task specific tools are actually more productive but on a more narrow area.

There exist universal algorithms for specific subareas such as learning. In theory, not working in practise.

Several subareas of AI have good theoretical foundations. No true intelligence has it so far.

6 Conclusion

AI systems can work well under favourable conditions, and are neither panaceas nor research curiosities. AI is not (just) art or a fashion, it is first of all a scientific discipline. At present, AI can importantly improve productivity and enhance the application areas of computers. As all other technologies, it must be used with a certain precaution and especially when circumstances are favourable. Therefore, more knowledge about AI in general as well as knowing about common legends and myths about AI may improve the success rate and extend the number of AI applications.

References

[1] Brooks R. A. (1991) Intelligence Without Reason. *Proceedings of the IJCAI'91 Conference*, Sydney, Australia, p. 569-595.

[2] Epstein R. (1992) Can Machines Think; The Quest for the Thinking Computer. AI Magazine, 13, 2, p. 80-95.

[3] Fox M. S. (1990) AI and Expert System Myths, Legends, and Facts. *IEEE Expert*, February 1990, p. 8-19.

[4] Gams M. (1992) Facts, Legends and Myths about AI. *Proceedings of the ERK'92*, Portorož, Slovenia, p. 139-142.

[5] Gams M., Drobnič M. & Petkovšek M. (1991) Learning from examples – a uniform view. *Int. Journal for Man-machine Studies*, 34, p. 49-89.

[6] Hayes P. J., Novak G. S. & Lehnert W. G. (1992) ACM Forum - In Defence of Artificial Intelligence. it Communications of the ACM, 35, 12, p. 13-14.

[7] Hayes-Roth F. & Fikes R. (1991) Interview. *IEEE Expert*, 6, 5, p. 3-14.

[8] Henery R. & Taylor C. (1992) Applications of Machine Learning, Statistics and Neural Networks in Prediction and Classification. *Proceedings of the ISSEK Workshop'92*, Bled, Slovenia, p. 22.

[9] Mettrey W. (1992) Expert Systems and Tools: Myths and Realities. *IEEE Expert*, 7, 1, p. 4-12.

[10] Michalski R. S. & Tecuci G. (1991) *Proceedings of the First International Workshop on Multistrategy Learning.* Harpers Ferry, USA.

[11] Minsky M. (1987) *The Society of Mind.* New York: Simon and Schuster.

[12] Schank R. C. (1991) Where's the AI?. *AI magazine*, 12,4, p. 38-49.

[13] Searle J. R. (1982) The Chinese Room Revisited. *Behavioral and Brain Sciences*, 8, p. 345-348.

[14] Sluga A., Butala P., Lavrač N. & Gams M. (1988) An attempt to implement expert system techniques in CAPP. *Robotics & Computer-Integrated Manufacturing*, 4, 1/2, p. 77-82.

[15] Urbančič T. & Križman V. (1991) *A Handbook of AI Applications.* IJS Press, Slovenia.

[16] Wilkes M. W. (1992) Artificial Intelligence as the Year 2000 Approaches. *Communications of the ACM*, 35, 8, p. 17-20.