

Expert systems and developing expertise: Implications of Artificial Intelligence for Education

Cite: Bruce, Bertram C. (1984). Expert systems and developing expertise: Implications of artificial intelligence for education. In Ralph Van Dusseldorp (Ed.), Alaska's challenge: Computers in education: Strategies and solutions (pp. 15-24). Anchorage, AK: Alaska Association for Computers in Education. Also as BBN Report No. 5732 (1984, June). Cambridge, MA: Bolt Beranek and Newman.

Few of those touched by the new world of computers have ever imagined that someday a powerful computer might step out of the realm of checkbook balancing and into that of Star Wars; that it could talk, listen and understand, reason, think and act intelligently. That dream, or nightmare, is the basis of Artificial Intelligence (AI), the field that studies the process of thought in order to produce machines that can do what most of us say requires true intelligence (see Rich, 1983; Winston, 1977)

AI is a new field, which draws from psychology, linguistics, mathematics, and philosophy. Each of these fields adds to our knowledge of how intelligent systems are constructed and what abilities they must have. Ultimately, these efforts can be seen as attempts to further our understanding of understanding itself.

This paper discusses a few issues in AI research with the aim of understanding whether the concepts or the tools of AI can be of use in education (see also Green, 1984). Most of the discussion focuses on natural language understanding, one aspect of the highly diverse field of AI.

Turing's Test

The question of whether computers could ever be truly intelligent must remain open for now. Part of the answer is that a computer can think if it can understand and communicate in a natural language. We would say that a computer understands a language to the extent that it recognizes the context and responds appropriately to sentences uttered in that context.

Along this line A.M. Turing (1963) proposed a test for intelligence based on his "imitation game." Player A tries to convince a judge that she (or he), and not Player B, is Player B. Meanwhile, B is doing the same. Of course, some mechanism, such as teletype communication, must be used so that the judge's decision is not too easy. Turing's operational definition of intelligence is that if Player A is a computer and if it does as well as a person in that role, then we should grant that the computer is intelligent. Passing the Turing test would certainly require a recognition of context and the ability to respond appropriately. (No current AI program could get past the first sentence in Turing's test.)

Language and Information Processing

Turing's test is a philosophical thought experiment that relates to a more general issue about thought and language: Before we can say that we comprehend the operation of an information processing system we must have an understanding of the way it encodes and transmits information, the symbols it uses and the way they are interpreted. In other words, we must have analyzed its language. This is true whether we are studying computer programs or the human mind. Language is so essential to the operation of an intelligent system that the study of language is almost inseparable from the study of thought, knowledge, perception, or problem solving. Postman and Weingartner's (1969) term "languaging" emphasizes the fact that in any discipline one essential task is creating a language, a system of symbols for structuring a particular view of the world.

AI research has been intimately connected with language research in two important ways: First, natural language understanding and generation have been major topics of study. Every AI system has needed some input/output language and natural language has typically been the desired one. Second, AI systems have themselves been developed using advanced programming languages. In fact, one of the major spin-offs of AI research has been the new languages created in response to demands of AI researchers or as manifestations of AI discoveries.

Using Context Appropriately

One of the principal issues in AI is how to use context to solve problems. Computer systems are impressive when given a well-defined task, but unlike humans they founder on novel problems or in new circumstances. They also fail to use those cues from the surroundings that any child would know how to utilize. A favorite domain in which to examine the use of context is natural language understanding (see Rubin, 1973; Winograd, 1972).

Natural languages include languages such as English, Latin, Russian, or Swahili but not formal languages such as BASIC, LOGO, or formal logic. Most of the plausible criteria for naturalness in a language, such as the presence of ambiguity, changing over time, extensibility, and use by humans, also now apply to many formal languages. Rather than making a sharp consider this: A language is natural to the extent that the meaning of a sentence in the language is a function of the "pragmatic context" in which the sentence is uttered. The pragmatic context includes characteristics of the surrounding discourse, the speaker, the hearer, the time and place, and the purpose of the conversation.

Abelson's Drugstores

Pragmatics is a broad and important aspect of language study, if not the whole of it. But there are less encompassing views of language which have been studied more thoroughly--namely, syntax and semantics. C. W. Morris gives roughly the following

definitions for these concepts: Syntax is the study of permissible orderings of symbols in the language (i.e., relations among linguistic expressions). Semantics includes syntax and also the study of the relationship between the symbols and the objects to which they refer. Pragmatics includes syntax and semantics and also the relationship of the language to its use and users.

Current natural language processors do very well on syntax; they do less well on semantics; and they almost always respond inappropriately whenever the pragmatic context changes. R. Abelson gives the following example which illustrates the concepts of syntax, semantics and pragmatics and also gives a brief but fair picture of where we are today:

Suppose a person says to a computer system " I went to three drugstores today." A syntactically based system, such as the first question answering programs, might respond "How did you go to three drugstores?" A semantically based system, knowing that a person goes to drugstores to buy things, might respond "What useful things did you buy in three drugstores?" A pragmatically based system might go further and say "Why didn't the first two drugstores didn't have what you wanted?"

National Language Understanding

Systems with rudimentary natural language understanding capabilities are being developed for information retrieval systems, computer aided design projects, robotics and medical history taking, record keeping and diagnosis systems. Computer assisted instruction systems are being designed which will allow natural language communication between a student and a computer. Natural language capabilities are essential for fully automatic programming and in simulations of cognitive processes.

The appeal of natural language understanding lies not so much in any one application as in the close connections now seen among language, knowledge, and problem-solving. In the small systems of today, we need a strong imagination to see much intelligence operating. Yet certain principles are emerging and among them the most important is that our understandings of language, of knowledge, and of problem-solving are intimately related. One hope is that as we learn more about natural language, especially in less restricted contexts, we will be learning more about the human mind as well.

Getting a computer to go beyond text to understanding the spoken word is one of the long range goals of work in natural language understanding (Bruce, 1981). It is much more difficult than understanding written text and in fact requires all the syntactic, semantic and pragmatic abilities of a written language system, plus extra abilities for speech. Aside from its intrinsic importance, speech understanding is interesting because it seems to require, more strongly than other applications, the integration of a system's various components.

Advances In Artificial Intelligence

AI is a rapidly moving field. Often the best work of a decade is obsolete within five years. Part of the reason for this rapid progress is the development of better tools, such as new programming languages and computer systems. In turn, many new programming language features have arisen through AI research. For example, the need to store and manipulate complex data structures led to the development of list processing languages such as LISP (a widely used AI language and forerunner of LOGO). The need for string manipulation capabilities and pattern matching led to languages such as COMIT, SNOBOL, and METEOR. The usefulness of recursive control structures was seen early and embodied in LISP and other languages. Later, as multiprocessing or parallel processing became desirable languages with the ability to back-track, such as GOL and PLANNER were developed.

Better languages have not been the sole explanation for improvements in AI systems. There have also been changes in representations for knowledge and in overall design.. The earliest approaches included simple pattern matching, hierarchical data structures, and highly specific deductive capabilities. In general, early AI systems tended to use specialized representations which restricted the domain in which they could operate.

Impact of AI on Education

A field such as AI, which examines the processes of learning, reasoning and language use and which pushes the limits of interactive computer technology might be expected to contribute significantly to education. There are four areas in which this is taking place.

First, AI is an important new science that can be expected to increase in importance. The excitement and relative accessibility of AI ideas and products can stimulate students to learn more about computer science and science in general. More fundamentally, students need to understand something of the basic concepts of AI because these concepts are essential in the development and continual reformulation of modern world views. They need to know about both the underlying science and the potentials and limitations, the good and bad effects of AI as technology.

Second, to the extent that AI contributes to our knowledge about thinking, learning, using language, and recognizing patterns, it may help us to understand how children can be helped in their learning. Already, the AI approach has had a major influence on psychology, linguistics and philosophy, three of the fields which spawned AI. Theories in reading and math education have likewise been influenced by AI.

Third, AI is at the leading edge of developing computer technology. The emphasis in AI has always been on making computers that are more sensitive to context, more able to recover from errors, more creative, in short, more like people. This has resulted in

tools--graphics, programming languages, debugging systems, and so on--that have proved useful far beyond the studies within AI. Many of these tools are now making their way into schools.

Finally, there are embodiments of AI theories about learning in tools designed specifically to teach. For example, J. S. Brown and R. Burton (see Sleeman and Brown, 1982), developed a computer coach for the arithmetic game, "How the West Was Wwon." The coach had to be an expert in playing the game but also had to be able to monitor a student's performance and to develop a model of what that student knew or didn't know. It then had to use its knowledge of arithmetic, the game, and the student to provide useful help.

Conclusion

AI is a topic with many faces. It is a scientific field of study that draws from the physical and social sciences, mathematics and the humanities. It is also a branch of technology that encompasses both computers and the interaction of people with machines. It should not surprise us to see AI begin to affect education. The important question is whether we can understand the potential and the limitations of AI well enough to channel its impact in beneficial ways.

References

- Bruce, B. Conversing with HWIM. Cognition and Brain Theory, 1981, @u(4), 123-138.
- Green, J. O. Artificial intelligence and the future classroom. Classroom Computer Learning, January 1984, pp. 26-31.
- Postman, N. & Weingartner, C. Teaching as a Subversive Activity. New York: Dell, 1969.
- Rich, E. Artificial Intelligence. New York: McGraw-Hill, 1983.
- Rubin, Andee, "Grammar for the People: Flowcharts of SHRDLU's Grammar", MIT AI Lab., AIM-282, 1973.
- Sleeman, D. & Brown, J. S. Intelligent Tutoring Systems. New York: Academic Press, 1982.
- Turing, A.M. "Computing machinery and intelligence." In E. A. Feigenbaum and J. Feldman (Eds.), Computers and Thought. New York: McGraw-Hill, 1963.
- Winograd, T. Understanding Natural Language. New York: Academic Press, 1972.
- Winston, P. Artificial Intelligence. Reading, Mass: Addison-Wesley, 1977.