the University of Alberta's ExpressNews your online news source

To print this story, click the "**Print**" button on your browser's toolbar

What is Artificial Intelligence?
By Dr. Russ Greiner


Dr. Russ Greiner

Email this story

With the excitement about the recent Spielberg film "A.I.: Artificial Intelligence", many people have started wondering about this field, asking what exactly is AI? As *intelligence* appears to be a distinctly human trait, many assume that AI is about artificial *human* intelligence; that is, producing machines that emulate human behaviour. The Spielberg film goes beyond this to suggest that AI involves human-like emotions as well.

Not all artificial intelligensia share this view. The prevailing view holds that there are many "species" of intelligence. Just as jets are not constrained to fly by flapping their wings, similarly, computers are not constrained to "think" in the same ways as people. To illustrate this point, consider the University of Alberta Chinook program, which in 1994 won the world championship of checkers over all comers, both man and machine.

To achieve this super-human level of performance, the Chinook authors (Prof Schaeffer and colleagues) tailored their program to exploit the strengths of the computer (fast processors, repetitive computations, large infallible memories), avoiding human-like reasoning processes that are hard to implement on a computer (reasoning by analogy, case-based reasoning, etc). The result was a program that could analyze an average of 21 moves into the future when deciding which move to make. This required considering 200,000 board positions per second; few humans can match that kind of processing power! In addition, the program used 444 billion pre-computed positions as part of its knowledge (all positions with eight or fewer pieces on the board).

No human can match that size of memory. Clearly, the computer's methods for achieving intelligence (here world-class checkers play) has some significant advantages over the human's model of intelligence. Indeed, one of the major scientific advances of the 20th century is the realization that we can create intelligence--in some cases, at a super-human level--using non-human entities.

There are a number of other AI projects at the University of Alberta that similarly achieve high-level performance. One, which has implications for Web search engines, deal with *question answering*; that is, finding answers to queries--for example, answering the question, which company manufactures Vauxhall?--from a large collection of documents, such as the entire World-Wide Web. Computers can use their speed and storage capacities to deal with billions of documents, which is beyond the scope of even the most gifted research librarians. However, while a person would immediately see that the sentence 'General Motors's Vauxhall car plant in Luton will close, costing about 2,500 jobs,' answers the query posed above, simple computer search algorithms would bypass it, as that sentence does not include the word 'manufacture'.

Prof Lin and his colleagues are addressing this limitation. Their system will automatically learn, from text, that the expression *"X manufactures Y"* is related to *"X's Y plant"*. These paraphrasal expressions can then be used to provide broader coverage and higher accuracy for question answering systems. Once incorporated, this will help reduce our frustration in the poor performance of search engines.

We can view this work as *learning* the relevant patterns from a body of examples. There are also a number of other tasks explicitly related to "computer learning". For example, Profs Greiner and Holte and their colleagues have worked on systems that can be attached to a computer program to automatically customize the computer's interface to each individual user.

These systems learn about the user by watching what s/he types in the normal interaction with the program. They detect patterns in the interaction and use them in a variety of ways to make using the program faster and more convenient for the user. To illustrate the general idea, consider standard mailers, which allow a user to store individual e-messages into folders created by the user. A learning system could watch the user and analyze the emails that were stored into each folder, and eventually learn user-specific rules that would predict which folder is appropriate for each email. It could then suggest appropriate folders for subsequent e-messages.

As each individual user has his/her own specific folders and dispatch rules, this could not be built into the software when it was developed; instead, it can only be determined after observing each individual user. This same idea--adapting a system by watching a user--can be used in video games, to help the game "learn" the patterns of the human player, either to help produce a "worthy adversary", or to produce "cloned player agents" that act like the human. Other "learning projects" attempt to reconfigure a Web site (first learn where each class of users typically want to go, then help these users find their destination pages more efficiently), learn patterns in DNA sequences towards interpreting genomic data, and learn which landmarks a robot should use when navigating, etc.

In addition to this work in games, query answering and learning, there is also a large body of work related to datamining (eg, content-based image retrieval, patterns in satellite data), cognitive agents (eg, to help software agents better communicate with humans), robotics (eg, to play real-time soccer) and vision (eg, to count trees from aerial photographs of forests, over all of Canada); see University of Alberta AI webpages.

Each of these systems was designed to help solve real-world problems. There is also a large body of more theoretical work, trying to better understand the fundamental principles of intelligence--eg, why are certain puzzles hard, but most day-to-day reasoning is easy; what concepts can or cannot be learned from what types of data; what are appropriate ways of (computer) agents to communicate (should they ever lie to each other?), etc.

These results can help us design more effective AI systems. While the U of A has become a prominent member of the AI community (in August 2002 we will host many of the world's largest annual AI conferences) there are of course a number of AI groups at other centres around the world, working on a large number of issues.

Among other accomplishments, there are now AI systems that diagnose various types of diseases, drive cars, recognize handwriting and spoken speech, evaluate credit card applications, translate text, control factories, schedule tasks, guide image-based robotic surgery, and troubleshoot complicated machinery.

As this list shows, AI is already pervasive, even though many people may not have realized it yet. There is a common theme to these impressive accomplishments:
AI involves producing programs that behave effectively in complex environments--this typically requires reasoning, possibly in ways that differ from how people think.
Doing this has required significant advances in understanding intelligence, both in the abstract, and in the interest in solving a wide variety of specific, important tasks. AI continues to be one of the most exciting, and useful, of fields today--with an impressive history of major results, and an "application pull" motivating yet other new advances, which will continue to improve the quality of our lives.

For more information, we encourage readers to explore AITopics website (which also includes a discussion of the recent Spielberg movie) or see the recent textbooks
? S Russell and P Norvig, Artificial Intelligence: A Modern Approach,

Prentice Hall, 1995.
? D Poole, A Mackworth and R Goebel, Computational Intelligence: A Logical Approach, Oxford, 1998.