Artificial Intelligence: Hype or Reality?

Adrian A. Hopgood
Nottingham Trent University

AI has been a rich branch of research for more than 45 years. Any sense that the field is overhyped probably results from a failure to appreciate the incredible complexity of everyday human behavior.

Most definitions of artificial intelligence in the standard texts are overly complex for a general survey of the field, so here is a simple one that will suffice instead: Artificial intelligence is the science of mimicking human mental faculties in a computer.

Pinpointing the beginning of AI research is tricky. George Boole (1815-1864) had plenty of ideas on the mathematical analysis of thought processes, and the field still retains several of his ideas. However, since he had no computer, my simple definition rules out Boole as AI's founder.

Just as historians on either side of the Atlantic have different opinions about who built the first programmable computer, they also diverge over AI's origins. British historians point to Alan Turing's 1950 article that defined what is now known as the Turing test for determining whether a computer displays intelligence.[1] American historians prefer to point to the Dartmouth conference of 1956, which was explicitly billed as a study of AI and is believed to have been the source of the term "artificial intelligence."

A SPECTRUM OF INTELLIGENT BEHAVIOR

A difficulty with my simple definition of AI is that it leaves the notion of intelligence rather vague. Figure 1 helps clarify the matter by showing a spectrum of intelligent behaviors based on the level of understanding involved. The lowest-level behaviors include instinctive reactions, such as withdrawing a hand from a hot object or dodging a projectile. High-level behaviors demand specialist expertise such as in the legal requirements of company takeovers or the interpretation of mass spectrograms.

Researchers have developed conventional computing techniques to handle the low-level decision making and control needed for the low end of the spectrum. Highly effective computer systems exist for monitoring and controlling a variety of equipment. For example, Figure 2 shows the Advanced Step in Innovative Mobility robot (http://asimo.honda.com). Developed by Honda, Asimo has 16 flexible joints, requiring a four-processor computer just to control its balance and movement. Owing to its rigid spine, the robot needs slightly elongated arms to pick up objects from the floor. Asimo shows exceptional human-like mobility, but it cannot think for itself. Despite the addition of voice and gestures, its behavior is still firmly anchored at the lower end of the spectrum.

Published by the IEEE Computer Society

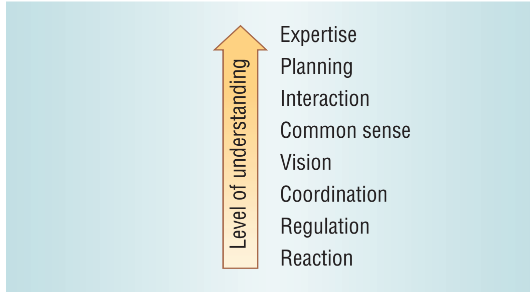

Figure 1. A spectrum of intelligent behavior, ranging from reaction-based to specialized expertise.

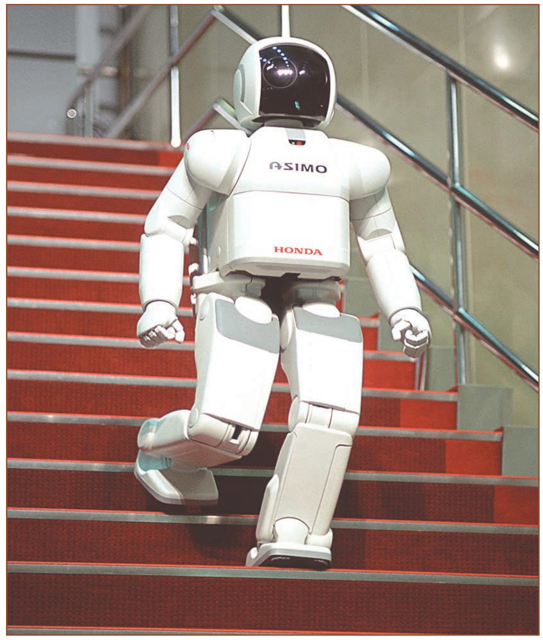

Figure 2. Honda's Asimo robot descending a flight of stairs. Asimo has 16 flexible joints with a four-processor computer to control its balance and movement, but its intelligence is limited. (Source: Honda. Used with permission.)

Figure 3. Spot the rabbit. Deciphering a photograph is a complex behavior that humans perform instantly but computers perform with difficulty, if at all.

Early AI research, on the other hand, focused on problems at the high end of the spectrum. Two applications, for example, concerned the specialist areas of mass spectrometry[2] and bacterial blood infections.[3] These early triumphs generated great optimism. If a computer could deal with problems too difficult for most ordinary people, then surely more modest human reasoning would be straightforward.

Unfortunately, this is not so. Human behavior in the middle of the spectrum, which we perform with barely a conscious thought, has proved to be the most difficult for computer scientists to emulate. Although computer software such as Mathematica can perform advanced calculus, a computer still cannot reliably recognize objects in a visual image.

Consider the photograph in Figure 3. Most of us can spot the rabbit in the photograph (though only a privileged few know that her name is Fluffy). But this perception is a complex behavior. First, recognizing the boundary between objects is difficult. Even after an object is delineated, recognition is far from straightforward. For instance, rabbits come in different shapes, sizes, and colors. They can assume different postures, and they may be partially occluded, as in Figure 3. A fully sighted human can process these perceptions in an instant, without considering it a particular mark of intelligence. But getting a computer to decipher a photograph reveals the task's astonishing complexity.

PRECONCEPTIONS FROM POPULAR FICTIONS

Fictional characters such as Hal in *2001: A Space Odyssey*, Robbie in *The Forbidden Planet*, and, more recently, David in Stephen Spielberg's film *AI* further fuel the huge expectations that the term "artificial intelligence" conjures up. In each case, the fictional intelligent system interacts with its environment through communication and action. In the latter two examples, it also resides in an android body.

Current technology lags behind these popular images, making it easy to find an audience for the idea that AI has failed to live up expectations. Yet, we clearly cannot use these fictions as a standard for measuring progress in the field. If AI were named "nifty computer programs," it would surely be hailed an unqualified success.

TWO APPROACHES TO AI

Current technologies for AI fall into two broad categories:

- explicit modeling with words and symbols, and
- implicit modeling with numerical techniques.

The first category includes techniques such as rule-, model-, frame-, and case-based reasoning. In this category, explicit rules may model the problem:

> If the pressure is high and the release valve is closed, then the release valve is stuck.

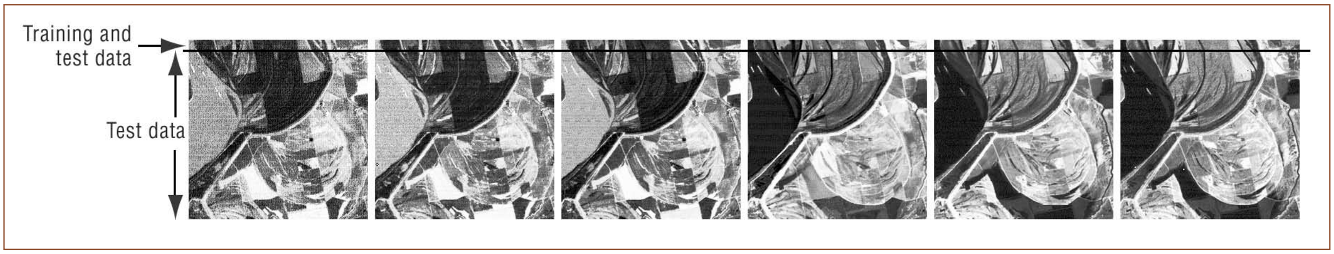

Figure 4. Portion of a Landsat-4 satellite image taken in six different wavebands. (Source: NASA. Used with permission.)

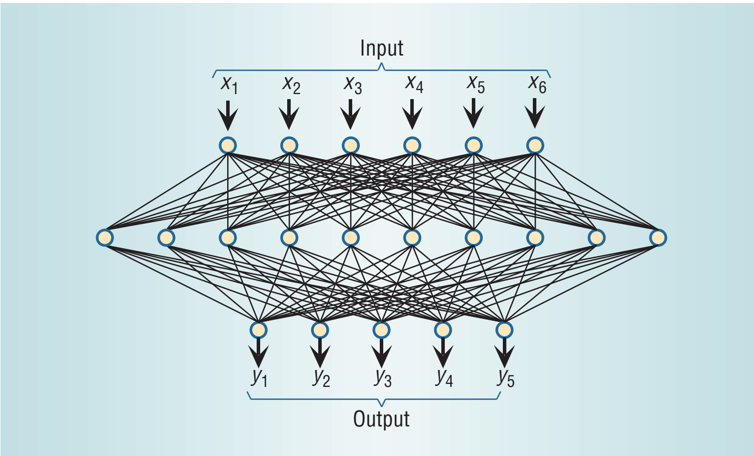

Figure 5. A multilayered perceptron. The three-layer network takes input data from six wavebands and outputs images according to five categories of land use.

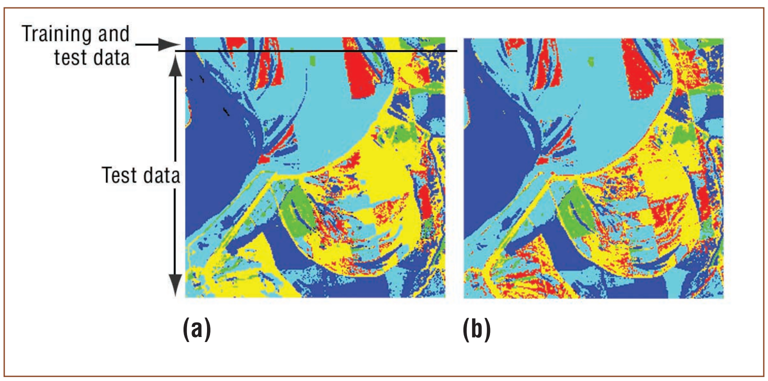

Figure 6. (a) Actual land use map and (b) land use map from neural network outputs. Key: dark blue = water; pale blue = swamp; green = trees; red = cultivated land; yellow = soil/rock; black = unknown.

More sophisticated variants attempt to take account of uncertainty:

> If the pressure is high and the release valve is closed, then the release valve is *probably* stuck.

Although these techniques have succeeded in their narrow domains, they are intrinsically limited. They handle only explicitly modeled situations and cannot deal with the unfamiliar.

The numerical approaches overcome these difficulties to some extent by enabling the computer to build up its own model, based on observations and experience. Neural networks, in particular, can learn associations from a set of examples and then apply them in previously unseen cases. These associations are particularly effective at classifying data patterns.

For example, Figure 4 shows a region of the Mississippi Delta, imaged at six different wavebands from a satellite. Figure 5 shows a simple neural network architecture, the multilayered perceptron, which was trained to associate the six waveband images with the corresponding land use. The image pixels constitute the inputs, and five land-use categories constitute the outputs: water, trees, cultivated land, rock, and swamp. The network was trained pixel-by-pixel on just the top 1/16th of these images and tested against the whole images.

Figure 6 shows the results. The classification is mostly correct, although there are some differences between the results and the actual land use. We could refine the neural network's parameters to further improve its performance, but even with the apparent discrepancies, Figure 6 clearly demonstrates the network's ability to generalize from a limited set of examples.

Interest in neural networks surged in 1985, following the discovery of an effective learning algorithm.[4] However, neural networks also have fallen victim to hype, perhaps because their name conjures up unrealistic notions of an artificial brain.

Other techniques such as genetic algorithms, simulated annealing, artificial immune systems, and fuzzy logic are useful.[5] They all come under the general heading of AI, but none really exhibits intelligent behavior.

BLACKBOARD SYSTEMS: INTEGRATING DIFFERENT TECHNIQUES

Blackboard approaches to building AI systems have proved useful in a variety of applications, regardless of whether the resultant system can be accurately described as intelligent.

The blackboard model is analogous to a team of experts gathered around a blackboard, on which a problem or some data appears. Any team member can contribute to solving the problem or interpreting the data, and a solution emerges on the blackboard. A blackboard system uses an area of shared computer memory to replace the physical blackboard, and the experts are different software modules.

Figure 7. DARBS application. The system interacts with plasma deposition equipment through sensors and actuators. (Photo by Lars Nolle.)

The blackboard approach recognizes that different subtasks require different techniques. The first published application of a blackboard system was Hearsay-II for speech understanding in 1975.[6] Further developments took place during the 1980s,[7] and the blackboard model is now seen as a key technology for the burgeoning area of multiagent systems.[8] In the late 1980s, my colleagues and I at the Open University developed the Algorithmic and Rule-Based Blackboard System, which we subsequently applied to diverse problems including the interpretation of ultrasonic images[9] and the control of plasma deposition processes.[10]

More recently, ARBS was redesigned as a distributed system, DARBS, in which the software modules run in parallel, possibly on separate computers connected via the Internet.[11]

SITUATED AI: INTERACTING WITH ENVIRONMENTS

Our most recent application of DARBS extends our work on controlling plasma deposition processes. As Figure 7 shows, the application is *situated*—that is, like the fictional intelligent computers, it interacts with its environment through sensors that detect the environment and actuators that operate upon it.

The Trilobite autonomous vacuum cleaner (http://trilobite.electrolux.se) illustrates recent progress in situated AI. As Figure 8a shows, the Trilobite is a commercially available robot developed in Sweden by Electrolux. The robot borrows the shape and name of the prehistoric animal shown in Figure 8b that cleaned the bottoms of the oceans 250-560 million years ago.

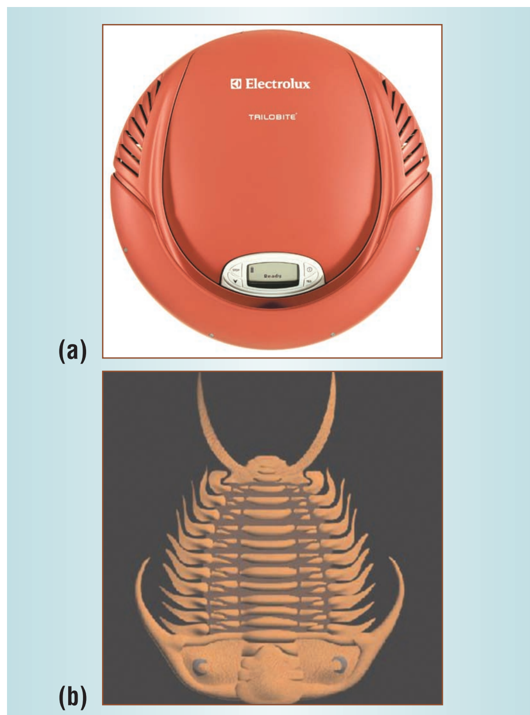

Figure 8. (a) A Trilobite autonomous vacuum cleaner from Electrolux; (b) a fossil trilobite—a primitive animal that cleaned the ocean bottom 250-560 million years ago.

The Trilobite vacuum cleaner embodies the principles of an AI planning system:

- it builds a model of its environment by circumnavigating the walls of a room;
- it establishes a goal—specifically, to traverse the entire accessible floor area;
- it assembles a set of actions to achieve the goal; and
- it replans as required to cope with changes in the world model, such as obstacles in its path.

Figure 9 shows Kismet (http://www.ai.mit.edu/projects/humanoid-robotics-group/kismet/), a robot head developed at the Massachusetts Institute of Technology. Kismet has no body or limbs, but it can communicate through hearing, vision, speech, facial expression, and posture. Described by its makers as a "sociable robot," Kismet uses nine

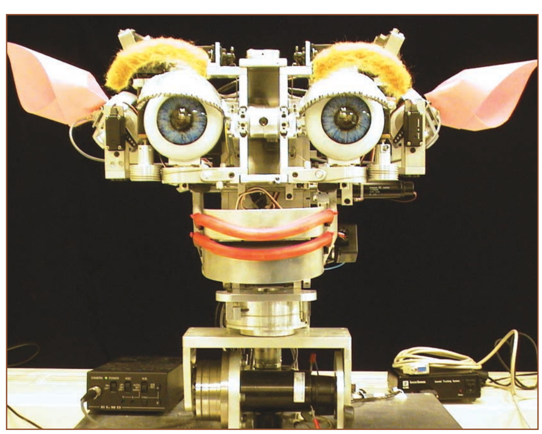

Figure 9. MIT's Kismet. The sociable robot communicates through hearing, vision, speech, and facial expression. (Source: Cynthia Breazeal, MIT Artificial Intelligence Laboratory. Used with permission.)

computers for vision alone and has 21 motors for controlling its face and neck. Investigators use Kismet to study the social interaction between the robot and the people it meets, especially children.[12]

Artificial intelligence has made significant advances from both ends of the intelligence spectrum shown in Figure 1, but a gap still exists in the "common sense" region. Building a system that can make sensible decisions about unfamiliar situations in everyday, nonspecialist domains remains difficult.

Nevertheless, it seems likely that applications like distributed blackboards will continue apace in specialist domains. This quiet revolution that attracts little press attention can make a significant contribution to our lives. It includes the current trend toward embedding AI within other hardware and software systems.

Meanwhile, we can also expect progress toward the popular idea of AI manifested in human-like robots. This development requires progress in simulating behaviors we tend to take for granted—specifically vision, language, common sense, and adaptability.

AI is not just hype, as many real working applications use AI techniques. However, AI is not yet a reality either, as we have yet to develop a system that spans the full spectrum of intelligent behavior. ■

Acknowledgments

The following people contributed to the development of ARBS, DARBS, and their applications: Lars Nolle, Patrick Wong, Kum Wah Choy, Gary Li, Heather Phillips, Hua Meng, Neil Woodcock, and Nicholas Hallam. Nick Braithwaite and Phil Picton cosupervised the plasma control project. The demonstration of a neural network for interpreting satellite images is based on a project originally devised in collaboration with Tony Hirst and Hywel Evans, using data supplied by Paul Mather. This article is based on my inaugural professorial lecture at the Nottingham Trent University.

References

1. A.M. Turing, "Computing Machinery and Intelligence," *Mind*, vol. 59, 1950, pp. 433-460.
2. B.G. Buchanan, G. Sutherland, and E. Geirgenbaum, "Heuristic DENDRAL: A Program for Generating Explanatory Hypotheses in Organic Chemistry," *Machine Intelligence 5*, B. Meltzer and D. Michie, eds., Elsevier, 1969.
3. E.H. Shortliffe, *Computer-Based Medical Consultations: MYCIN*, Elsevier, 1976.
4. D.E. Rumelhart, G.E. Hinton, and R.J. Williams, "Learning Representations by Back-Propagating Errors," *Nature*, vol. 323, 1986, pp. 533-536.
5. A.A. Hopgood, *Intelligent Systems for Engineers and Scientists*, 2nd ed., CRC Press, 2001.
6. V.R. Lesser et al., "Organization of the Hearsay-II Speech Understanding System," *IEEE Trans. Acoustics, Speech, and Signal Processing*, vol. 23, 1975, pp. 11-24.
7. H.P. Nii, "Blackboard Systems, Part One: The Blackboard Model of Problem Solving and the Evolution of Blackboard Architectures," *AI Magazine*, vol. 7, 1986, pp. 38-53.
8. G. Brzykcy et al., "Multiagent Blackboard Architecture for a Mobile Robot," *Proc. IEEE RSJ Int'l Conf. on Intelligent Robots and Systems*, vol. 4, IEEE Press, 2001, pp. 2369-2374.
9. A.A. Hopgood et al., "Interpreting Ultrasonic Images Using Rules, Algorithms, and Neural Networks," *European J. Non-Destructive Testing*, vol. 2, 1993, pp. 135-149.
10. A.A. Hopgood et al., "Fuzzy Logic in a Blackboard System for Controlling Plasma Deposition Processes," *Artificial Intelligence in Engineering*, vol. 12, 1998, pp 253-260.
11. L. Nolle, K.C.P. Wong, and A.A. Hopgood, "DARBS: A Distributed Blackboard System," *Research and Development in Intelligent Systems*, vol. 18, M. Bramer, F., Coenen, and A. Preece, eds., Springer, 2001, pp. 161-170.
12. C. Breazeal and L. Aryananda, "Recognition of Affective Communicative Intent in Robot-Directed Speech," *Autonomous Robots*, vol. 12, 2002, pp. 83-104.

Adrian A. Hopgood *is professor of computing and head of the School of Computing and Mathematics at the Nottingham Trent University, UK. He is also a visiting professor at the Open University. His main research interests are in intelligent systems and their practical applications. Hopgood received a PhD in materials physics from the University of Oxford in 1984. He is a member of the British Computer Society and a committee member for its specialist group on artificial intelligence. Contact him through his Web site at www.adrianhopgood.com.*