Artificial Intelligence: Can Machines Think?

by Raymond L. Paden and James Wolfer

Horses gave way to automobiles. Rubber to plastics. Cotton to polyester. What about intelligence? Are computers about to replace human intelligence? Is artificial intelligence a threat to our humanness?

The term artificial intelligence (AI) is rather new. Over this century, theories concerning AI co-evolved around two perspectives: a formal approach using deterministic computer programs and a biological approach.

Understanding Formal AI

AI immediately evokes images of user-friendly robots from movies like *Star Wars* or *Star Trek*. In reality, however, the meaning of AI is rather nebulous. Note how varied the definitions are:

- "AI is the attempt to answer the question ... how does the human brain give rise to thoughts, feelings, and consciousness."
- "AI is the study of computer problems that have not been solved."
- "AI is the art of creating machines that perform functions that require intelligence when performed by people."[1]

Each of these statements define AI in a manner suitable to the goals of particular research interests. But none of them defines AI conclusively. So what do scientists mean when they talk about formal AI? Basically, they refer to a deterministic computer program capable of imitating intelligent behavior.

Historical background. AI emerged from the fertile turmoil of mathematics from 1870 to 1930, when the goal was to unify all mathematics using a small collection of basic principles. However, this goal remained elusive. The most ambitious of these attempts was set forth by David Hilbert as a mathematical problem, known as the *Entscheidungsproblem* ("the tenth problem").

Hilbert's goal was to prove that mathematics is consistent (without contradictions), complete (all mathematical statements could be proved or disproved), and computable (the truth value of any mathematical statement can be determined by a mechanical device). However, to avoid the difficulties associated with other attempts to unify mathematics, the problems and their proofs were approached on strictly formal methods, that is, following logical rules of inference based on axioms. Such formal methods would substitute human insight and judgment with mechanical means.[2]

Hilbert's *Entscheidungsproblem* was ultimately proven to be impossible by the logician Kurt Goedel, who demonstrated in 1931 that mathematics could not be both complete and consistent. While Goedel ignored the issue of computability, his finding nonetheless modified the question associated with it to ask: "Does there exist an algorithm to decide if a problem has a solution?"[3] In 1936 a theoretical model of computation, the Turing machine, proved that even this was not possible.

The Turing machine (TM, see box, p. 11), developed by the mathematician Alan Turing, is a mechanical device that precisely defines the notion of an algorithm. In other words, it outlines the steps one must follow to accomplish a task—a "recipe," if you please. The machine is programmed to solve a problem defined in strictly formal terms. However, although such problems do not require semantic insight—that is, knowing the meaning behind a rule or a symbol—they can potentially represent

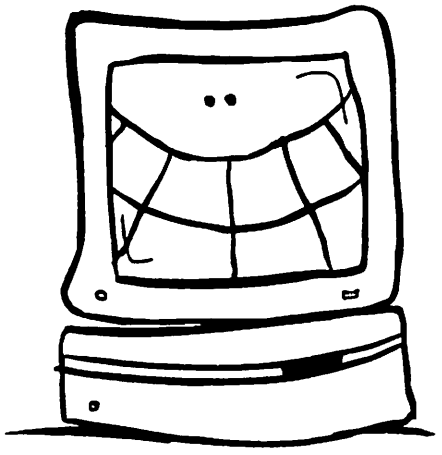

real-world problems ranging from balancing a bank account to simulating the behaviors of human intimacy. Today's computers are equivalent to a TM, in that any problem that can be solved by a TM can be solved by a computer and vice-versa. (Hence the terms *computer* and *TM* are used interchangeably in this essay.) Problems that can be solved, or that can be attempted but are not solvable by a TM, are called computable. Problems that are so complex that they cannot even be attempted on a TM are called uncomputable.

Church's Thesis and the Goal of Formal AI. A TM is a simple machine. In spite of this simplicity, it is believed to be the most powerful form of mechanical computation known to man. It can execute any procedure that can be mechanically performed. While this assertion regarding TM's power, known as Church's thesis, cannot be formally proven, no mechanical model of computation that has been invented is more powerful.[4] Going a step farther, we can ask another question: Can all of human intelligence be performed by mechanically equivalent procedures?

To this, Hofstadter responds rhetorically: "Here one runs up against a seeming paradox. Computers by their very nature are the most inflexible, desireless, rule-following of beasts. Fast though they may be, they are nonetheless the epitome of unconsciousness. How, then, can intelligent behavior be programmed? Isn't this the most blatant of contradiction of terms?"[5]

To the proponents of formal AI, including Hofstadter, this is not a contradiction at all. Indeed, they believe that the next century will have computers that are the functional equivalents of human beings. However, the critics of formal AI, like Joseph Weizenbaum, assert: "We are capable of listening with the third ear, of sensing living truth, that is truth beyond any standards of provability. It is that kind of understanding, and that kind of intelligence that is derived from it, which I claim is beyond the abilities of computers to simulate."[6]

Perhaps, it can be speculated, this living truth—which represents higher forms of human intelligence—lies in the domain of uncomputable functions that are not accessible to computers.

The meaning of humanness

The attempt to emulate human intelligence raises the question of our humanity. Christianity and science have disagreed over this issue for many centuries.

The Scripture looks at humans from the perspective of Creation and redemption. It raises the question, "What are human beings?" (Psalm 8:4, NRSV), and provides some answers. Humans are created "in the image of God" (Genesis 1:27). They are spiritual (Romans 8:16; 1 Corinthians 2:11, 14-16), intellectual (Isaiah 1:18, Mark 12:30), creative (Exodus 31:1-5; Psalm 33:3), social (Genesis 2:18), affectionate (Proverbs 18:24; Ecclesiastes 3:5) and sexual (Genesis 4:1; Song of Solomon 4:16-5:1). God has given human beings freedom of choice (Deuteronomy 30:19; Joshua 24:15; John 7:17), but this freedom is not absolute (Romans 6:23). God has made them to be loving creatures (Matthew 22:37-39), but they also have the capacity to hate (Ecclesiastes 3:8). Through their choice, they have fallen (Romans 5:12,17), but God has sent His Son (John 3:16, Philippians 2:6-11) to restore them into His image (Acts 3:21; 1 John 3:2), provided they consent (John 14:15). Moreover, God will hold them accountable for their choice in the judgment (Ecclesiastes 12:13, 14).

Science looks at human beings from a different perspective altogether. Formal AI views the mind in behaviorist terms, based on the doctrine of logical positivism.[7] The mind is conceived of as a machine, and the task of formal AI is to create another machine, a properly programmed computer to be the mind's equivalent.[8] Thus science typically ignores many of the issues that arise from the biblical perspective.

At first glance, these two views of humanness seem at odds; but are they? One must ask whether there is anything special about the mechanisms of the brain. Be it computer chips or biochemical processes, the issue is intelligence, not the hardware that sustains it. Moreover, proponents of formal AI would argue that either these lofty aspects of our humanity can be programmed or are mere illusions.

Freedom of choice. A good example of the complexities involved in such questions can be seen in the issue of non-determinism (i.e., freedom of choice). In formally emulating intelligent behavior, one uses a computer that is programmed deterministically. But it can be proven that deterministic and nondeterministic TMs are equivalent. Thus scientists conclude that our feeling of free will can be programmed, using deterministic techniques.

Hofstadter accounts for this "feeling" of free will: "It is irrelevant whether the system is running deterministically; what makes us call it a 'choice maker' is whether we can identify with a high-level description the process which takes place when the program runs. On a low... level, the program looks like any other (deterministic) program; on a high... level, qualities such as 'will,' 'intuition,' 'creativity,' and 'consciousness' can emerge."[9]

Thus AI proponents argue that at the low level of neurophysiology deterministic choices are made in the brain, similar to the way they are made in a TM, and at high levels of consciousness, people merely have the perception of free will. Hence the argument: A human being is only an automaton.

If this is the case, as Bible-believing Christians we need to make three observations. First, this view of humans runs counter to basic Seventh-day Adventist teachings.[10] Second, it contradicts the Scriptures, which assert that human beings must choose whom they are to follow. Since we will be held accountable for this choice in the

judgment (Ecclesiastes 12:13, 14), we must have the ability to choose our destiny. Third, it denies that the love principle is God's ideal for humanity (Matthew 22:37-39; 1 John 4:8). Since free choice is the infrastructure of love, then we must be free to be able to love God! These observations should lead to question, if not reject, the proposed equivalence between the human mind and the TM.

The ethics of AI

Another question needs to be raised: Is it desirable to create a "machine" identical to a human being? From a practical perspective, many would probably answer No. In creating a machine, it makes little sense for scientists to program it to make arithmetic errors, to get angry, or to lie. Moreover, it would make no sense to program a computer to anticipate the future, only to have it "dismantled" once it becomes obsolete, and its software is no longer transferable to a new generation of computers.

If it were possible to create formal AI, many scientists would probably develop a machine with an alien intelligence that is understandable and submissive to people, much like robots in science fiction movies. Such machines would be programmed to recognize speech and have vision. They would be given a socially pleasing and accessible means of interfacing with humans. And they would need to be granted rights appropriate for an intelligent agent.

However, such a vision of formal AI may not be necessary for a TM model of genuine intelligence. Rather than creating truly intelligent machines, programs which merely appear intelligent could be designed to emulate those aspects of the human mind that are programmable. In this sense, intelligence becomes a practical metaphor used in the design of programs. The less accessible aspects of intelligence such as free will and spiritual vitality would not, could not, be programmed.

Biologically inspired AI

While formal AI has attained some success in areas such as expert systems and game-playing strategies, it has essentially failed to achieve many day-to-day survival necessities such as vision, which even the simplest creatures perform easily. Recent theoretical advances have produced a renaissance of paradigms based upon biological metaphors. These include artificial neural networks, genetic algorithms, genetic programming, and artificial life. These approaches share the basic assumption that complex behavior can emerge from simple computations or processes.

Artificial neural networks, for instance, are based upon the concept that useful computation can be distributed throughout a system of very simple computational elements ("neurons") by encoding information at the points of connections between these elements. Scientists have developed procedures that allow interconnected networks of these metaphorical neurons to learn relationships by example. Applications of these artificial neural networks include learning to drive a vehicle by "watching" a human driver, cancer screening, and financial management.[11]

In theory, one can apply the genetic mechanisms for information transfer in nature, such as selection, mutation, and sexual reproduction. One can also apply genetic mechanisms to search for a set of solutions under which an organism would survive, operating within a given set of stimuli and stipulated conditions. Genetic programming is an example of evolutionary computation that actually evolves programs to solve particular problems.[12] Applications of genetic programming include creating aesthetically pleasing art, learning to balance an inverted pendulum ("broom balancing"), and automatic target image recognition.

Artificial life research attempts to abstract the characteristics of life and reproduce them in some computational form. Farmer and Belin identify some of these attributes: life as a pattern in space-time (e.g. most of our cells are replaced in our lifetime); self-reproduction; information storage of self-representation (for example, DNA); metabolism; ability to interact with the environment; interdependence of parts forming the organism; stability under perturbations and small changes; and ability of the lineage to evolve.[13]

Those involved in artificial life research recognize two different claims, the strong and the weak. The weak claim asserts that anything produced is a simulation that may explain certain

The Turing Test and artificial intelligence (AI)

One significant dilemma facing proponents of formal AI is how to recognize intelligence in a computer when it is manifested, since there is no adequate definition for intelligence. Alan Turing attempted to answer this question in 1950 with an operational view of AI, using what is now called the Turing Test.[17] A computer claimed by its designers to be intelligent and a person are hidden from the view of a panel of judges. The judges interview the computer and the person through a keyboard and a monitor to determine which is the computer and which is the person.

Suppose a judge asks the respondents to factor a 30-digit integer. The answer would be a quick matter for the computer, but quite tedious for the person. Thus it would be necessary to program the computer to slow down on mathematical responses and even make mistakes. It would also be necessary to program the computer to get angry, lie, and cheat as well as to emulate the more noble aspects of humanity such as appreciating the aesthetic appeal of a musical composition.

So suppose it were possible for a computer to appear genuinely human (for example, to think and feel). Does this necessarily imply that a real aspect of humanity is alive in the computer's circuitry? If a computer acts intelligent by means of a program, then to the proponents of formal AI, it is intelligent. To others, if there is no understanding of the meaning behind a rule or a symbol (semantic insight), then there is no intelligence: "Acting, no matter how skilful, is not the real thing."[18]

properties of life. The strong claim asserts that the computer programs will eventually achieve the state of actually being "alive." Will a machine ever be intelligent? Will a machine ever be "alive"? Are our concepts of intelligence and life so focused on biological forms that we will preclude anything else attaining that status by unspoken definition? These questions cannot be answered at this time, but there is still much to gain from their study. As Langton asserts: "Although AI has not yet achieved anything that even its most ardent supporters would call genuine machine intelligence, AI has completely changed the way in which scientists think about what it is to be 'intelligent', and has, therefore, made a major scientific contribution, even though it hasn't achieved its overall goal."[14]

Similarly, research in artificial life will force us to rethink what it means to be "alive." Farmer projects some possibilities: "With the advent of artificial life, we may be the first species to create its own successors. What will these successors be like? If we fail in our task as creators, they may indeed be cold and malevolent. However, if we succeed, they may be glorious, enlightened creatures that far surpass us in their intelligence and wisdom. It is quite possible that, when the conscious beings of the future look back on this era, we will be most noteworthy not in and of ourselves but rather for what we gave rise to. Artificial life is potentially the most beautiful creation of humanity. To shun artificial life without deeper consideration reflects a shallow anthropocentrism."[15]

A Christian response to AI

From the current state of research in AI to consciousness, it is a long journey—and one that has achieved only small, incremental steps toward its goal. Yet the assumption of many, and the goal of some, is that it is not only possible, but also inevitable.

While the authors of this essay have differing opinions about the potential for creating artificially intelligent agents, they agree that we should be cautious about ruling it out categorically. Though science is incapable of discovering the totality of truth,[16] nonetheless, many of its experimental discoveries have produced tangible benefits. Moreover, we must always recognize that our arguments may be incomplete or even wrong. For instance, recall such events as the Great Disappointment of the Millerite movement or statements asserting that people would never land on the moon because they are sinful and the moon had known no sin. If we base our beliefs on fears of the unknown, they are likely to be shattered, resulting in a crisis of faith.

So how should a Christian respond? The Bible does not seem to directly preclude artificial intelligence. What Scripture does provide, however, is a stable basis from which to evaluate the consequences of artificial intelligence. Even if machine intelligence surpasses human intelligence in certain areas, we as Christians need not lose any of our self-worth or identity. Many feel threatened by AI's potential encroachment upon their humanity. In an age when we are at times reduced to numbers and intimidated by computers, is not the attempt to make machines our equal the ultimate threat to humanity? The answer is No. Our humanity is rooted in our relationship with our Creator, and our ultimate destiny is well defined in Scripture. Regardless of the successes or failures in AI, we must remember that God has "fearfully and wonderfully made" us (Psalm 139:14), that He sent His Son to redeem us (John 3:16; 1 John 2:1-2), and that we are welcome before His throne (Hebrews 4:16). Nothing can separate us from the love of God (Romans 8:38-39). This alone, if nothing else, sets us apart from machines.

Raymond L. Paden (Ph.D., Illinois Institute of Technology) and James Wolfer (Ph.D., Illinois Institute of Technology) are, respectively, chairman and associate professor in the Computer Science Department of Andrews University, in Berrien Springs, Michigan, U.S.A.

Notes and references

1. R. Kurzweil, *The Age of Intelligent Machines* (Cambridge, Mas.: MIT Press, 1990), pp. 13-15.
2. See R. Penrose, *The Emperor's New Mind* (New York: Penguin Books, 1989), pp. 102-105.
3. D. I. A. Cohen, *Introduction to Computer Science*, rev. ed. (New York: John Wiley & Sons, 1991), p. 806.
4. See M. Minsky, *Computation: Finite and Infinite Machines* (Englewood Cliffs, N.J.: Prentice-Hall, 1967), p. 108.
5. D. R. Hofstadter and Escher Goedel, *Escher, Bach: An Eternal Golden Braid* (New York: Vintage Books, 1980), p. 26.
6. J. Weizenbaum, *Computer Power and Human Reason: From Judgement to Calculation* (San Francisco: W. H. Freeman and Co., 1976), p. 222. Note: We are not aware of Weizenbaum's religious commitments, but his reference to "living truth" strikes a resonant chord within the Christian heart.
7. See H. Smith, *Beyond the Post-Modern Mind* (New York: Crossroads Press, 1982), p. 82; Kurzweill, *op. cit.*, p. 35.
8. For an example of the equivalence between computers and mind, see "A Conversation with Einstein's Brain" in D. R. Hofstadter and D. C. Dennett, eds., *The Mind's I* (Harmondsworth, Middlx., England: Basic Books, Inc., 1981).
9. Hofstadter, *Escher, Bach...*, pp. 713, 714.
10. *Seventh-day Adventists Believe...* (Hagerstown, Maryland: Review and Herald Publ. Assn., 1988), pp. 78-96.
11. See D. E. Windrow and M. A. Lehr, "Neural Networks: Applications in Industry, Business and Science," *Comm. ACM* 37:3 (March 1994), pp. 103-105; and T. Kanade and L. E. Weiss, "New Technologies and Applications in Robotics," *Comm. ACM* 37:3 (March 1994), pp. 58-68.
12. See "Genetic evolution and co-evolution of computer programs," in Christopher Langton, ed., *Artificial Life II* (Reading, Mass.: Addison Wesley, 1992).
13. See J. D. Farmer and A. Belin, "Artificial Life: the Coming Evolution," in *Artificial Life II* (Reading, Mass.: Addison Wesley, 1992).
14. Quoted in Langton, *Artificial Life II.*
15. Farmer, quoted in Langton, *Artificial Life II.*
16. See Smith, *Op. cit.*, p. 134.
17. A. M. Turing, "Computing Machinery and Intelligence," *Mind*, 59:36 (1950), pp. 433-460.
18. See Penrose, *Op. cit.*, pp. 5-11.