Lutheran Mission Matters

Volume XXIV, No. 3 (Issue 49) November 2016

Intelligent Computers in a Christian Worldview

Gary Locklair

Abstract: Thinking Machines! The inflection given when speaking these two simple words can invoke fear, excitement, concern, wonder, skepticism, or hope. This article explores the question of intelligent computers from a Christian worldview perspective. Both the origin and purpose of artificial intelligence are reviewed with an emphasis on how the field should be viewed and shaped within a Christian perspective. Answers to two fundamental AI questions will be presented: Can/will computers be intelligent, and can/will computers be equivalent to human beings?

Many people already believe that computers are intelligent. It is a confusing situation since computers appear to be intelligent. Imagine you are ready to buy a new Z06 Corvette. You wonder out loud to your neighbor, "I wonder what my monthly payments will be after a sizable down payment, seeing that I only have to finance $100,000 at 3.5% over five years?" If your neighbor was able to immediately calculate the answer ($1,819), you would surely be impressed and might respond, "You're quite intelligent." Imagine that you are lost in a strange city and stop to ask for directions. If the person helping you immediately drew a map showing 18 turns and the mileage between each to reach your destination, you would be dumbfounded by the stranger's intelligence. Of course, computer systems do these things all the time. The reason computers appear intelligent is due to the application of information theory. The information is being produced, ultimately, by intelligence, that is, intelligence infused into the computer system by intelligent people. Computers possess some attributes of their human creators. Because human intelligence is "behind" the creation and use of a computer system, the system reflects some appearance of intelligence.

Dr. Gary Locklair is professor and chair of Computer Science at Concordia University Wisconsin where he has served for thirty years. Prior to that he worked in the computing industry for Hewlett-Packard Company. He holds the BS, MS, and PhD degrees in Computer Science and is listed in several editions of Who's Who along with being a member of Upsilon Pi Epsilon, the computer science honor society, and the Alpha Chi honor society. Locklair is also an active researcher in the creation model of origins and is a founding member of the Concordia University System Society of Creation. locklair@cuw.edu

People

People are the most important aspect of a computer system because today people provide the intelligence of the system. Hardware and software together form a tool created by human intelligence and put to practical application by human intelligence. Although computers appear intelligent, their current intelligence is a reflection of the human intelligence used in their creation and use. Can computers truly be intelligent on their own?

When users work with an application package, they are reaping the benefits of the intelligence supplied by the programmers, information technologists, and computer scientists that created the system. The information produced by the algorithms is a result of the intelligent acts of all these people, along with the user, who must supply useful data. Programming captures the human thought process. When an application package is created, the algorithms are a reflection of how the programmer would solve the problem.

What's unique about people vis-à-vis problem-solving? The unique attributes of people, such as intelligence and creativity, are difficult to quantify but easily seen and understood. Children are natural born problem solvers. Children naturally love to explore and understand, as demonstrated by their favorite question: "Why?" There are powerful heuristic qualities "built in" to human beings.

Why are people intelligent and creative? The answer is that people were created in God's image. An omniscient, omnipotent God formed people with attributes similar to His own, but to a lesser degree.

Why are people intelligent and creative? The answer is that people were created in God's image. An omniscient, omnipotent God formed people with attributes similar to His own, but to a lesser degree.

Will computers think and be intelligent as people are? This is the active research question in the field of Artificial Intelligence. In order to understand the field of AI, we need to investigate its origins.

It was not long after the first true computers were completed in the 1940s that computer scientists began asking the question: "Can a computer do other tasks besides number crunching?" In 1956 a seminal conference on AI was held at Dartmouth, where a number of computer scientists exchanged ideas and developed an informal research agenda. The organizers of the conference declared, "Every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."[1]

The initial successes of the 1960s seemed to confirm the organizers' statement. By the 1960s, there were computer systems that could play a game of checkers and even defeat human opponents. In the 1960s Joseph Weizenbaum created the ELIZA system, which attempted to mimic a Rogerian psychotherapist. Some "sessions" of interaction were so successful that they led their human patients to believe they were interacting with a human being and not a computer. Examples such as these led to high hopes for AI. In the 1970s there were many predictions of thinking machines "just around the corner." One prediction was for the emergence of "robot servants" by the next decade! By the 1980s, there was the realization that general intelligence was an extremely difficult problem. By this time, the field of AI had split into two distinct camps.

Weak AI

One group of AI researchers decided that the "general intelligence" problem was either intractable or else not practical. Instead of creating machines that were autonomously intelligent in general, researchers in weak AI focused on simulating intelligent behavior in more narrowly defined areas. The goal of weak AI research has been realized and today can be seen in expert systems, as one example.

An expert system is a software application package that simulates the intelligent behavior of a human expert in a specific, narrow field of knowledge. Expert systems exist for many diverse fields, including medical diagnosis. A medical diagnosis expert system will accept symptoms as input and then produce a diagnosis as output. In other words, it is performing the same task as a human health care provider in the narrow field of medical diagnosis. Because an expert system accepts input, performs processing, and produces output, it is similar to any other software system. There are two distinguishing characteristics of an expert system worth noting: First, an expert system has a knowledge base. Similar to a database, a knowledge base contains the set of many, pertinent facts about the job or task. Secondly, an expert system has an inference engine. The inference engine applies rules of logic to the problem in order to select relevant facts and "reason" about them. The expert system is able to generate conclusions based upon the contents of the knowledge base and the rules of the inference engine.

One purpose of an expert system is to assist people. Originally, medical diagnostic expert systems were created as an aid for health care professionals. Because most doctors have not encountered every known disease, an expert system can aid them in diagnosing conditions that they have not before encountered. The expert system provides expert assistance with greater productivity and efficiency. Another purpose of an expert system might be to replace people. While this sounds outlandish, it is, of course, part of the advancement of a technological society. Many people would not want their doctor replaced by an expert system; yet there are numerous tasks performed more productively and efficiently by computers every day. In the twenty-first century, no one is concerned because human elevator operators began to be displaced by automated systems in the early twentieth century. If an expert system can fly a commercial airliner with the same effectiveness as a human pilot, is it farfetched to predict that one day airplanes will not have pilots?

Currently, a number of companies are researching autonomous vehicles. Today there are cars that drive themselves without the "assistance" of a human being. Some believe that in a few years autonomous cars will be common on the highways of the United States. Would you ride in a driverless car? Would you want to drive around other driverless cars?

Given the current state of computer science, I suspect that few people would fly on an airplane controlled exclusively by an expert system. Given the realities of bugs and the relative inflexibility of algorithms, the system does not appear to be robust. A word processor's crashing and losing a document is one thing; the crashing of an expert system onboard a commercial airplane is quite another. Yet, the point of AI research is to "break out" of the algorithmic mold and create systems that can reason about even unexpected situations and react accordingly. In other words, AI hopes to replace algorithms with heuristics such that computer systems can think and reason as humans do.

Strong AI

While the proponents of weak AI focus on simulating intelligent behavior in specific areas, the proponents of strong AI strive to create computers that are autonomously intelligent. One objective of strong AI research is to create thinking machines via human equivalence.

Because algorithms are somewhat inflexible, software applications are somewhat inflexible. Current generation software is not very robust; that is, current software is not able to respond to unforeseen circumstances. Strong AI researchers hope to overcome this problem by creating systems that are autonomously intelligent. A weak AI expert system that cannot "adapt" to new situations is not the ideal candidate to substitute for a pilot on a commercial airliner! However, an intelligent computer system that could react as a human being might be a candidate to fly airplanes autonomously.

One objective of strong AI research is to create thinking machines via human equivalence.

Computer Intelligence

Can computers be intelligent? The key term demanding definition in this question is "intelligent." Defining the term "intelligent" is a difficult task. Common dictionary definitions of intelligence usually include statements such as the ability to learn, understand, or deal with new situations; the ability to apply knowledge to manipulate the environment; and the ability to think abstractly. One approach to defining intelligence is to provide a framework of associated activities. In a hierarchical fashion, intelligence is seen to include activities such as thinking, reasoning and understanding (higher order), learning (middle order), and remembering and computing (low order). With this hierarchy constructed, it is now a straightforward process to assess whether computers can be intelligent.

Computers currently are unequaled at performing the low order activities of intelligence. Computers can compute and remember in ways far superior to people. The fastest computers at the beginning of the twenty-first century operate in the range of 10 teraflops (trillions of operations per second). For example, one special IBM RS/6000 SP computer can perform 12 trillion multiplications in a single second. It is difficult for a human being to fathom this computational speed, much less attempt to match it! Computers can also remember vast quantities of data. The previously mentioned IBM RS/6000 SP has been connected to a storage system with a capacity of 160 terabytes. It is estimated that 160 trillion bytes is enough storage to encode the entire information content of the United States Library of Congress—twice! Some researchers believe that human beings do store every experience they have throughout their lifetime. If this is true, then the human storage capacity is very great. However, as we all know, even if we store much, we often have trouble with recall (especially during a test)! Even if a person could remember as much as a computer, the computer can still recall (locate) the requested data more reliably than a person.

To illustrate the importance of definitions, consider the mid-twentieth century view that computers were "giant brains." Computer scientist Edmund Berkeley's 1949 book, *Giant Brains or Machines That Think*, illustrates the fundamental role definitions play in problem solving. Berkeley wrote, "A machine [computer] can handle information; it can calculate, conclude, and choose; it can perform reasonable operations with information. A machine [computer], therefore, can think."[2] If the definition of thinking (or intelligence) is limited to the low order activities, as Berkeley's is, then a computer is indeed an intelligent, thinking entity. If one's definition of intelligence is computation and memory, then a computer is super-intelligent, truly a giant brain. However, intelligence is more than just the low order attributes. Therefore, twentieth century computers did not actually think.

The situation is tricky to analyze, however. Consider another statement of Berkeley: "When you and I add 12 and 8 and make 20, we are thinking." Does addition require thinking? It is true that human beings both think and add, but does addition require human level thinking? No. A calculator can add two numbers, yet a calculator is not a thinking entity. Berkeley assumed that addition required thinking, and thus a device that performs addition is thinking. However, there is a known algorithm for addition, and both Wilhelm Schickard and Blaise Pascal created mechanical devices to implement the addition algorithm in the seventeenth century. These mechanical devices were not thinking when they added; they were merely following the rules encoded in their mechanisms. The manner in which one defines terms is important for understanding and answering questions.

What about the middle order activities associated with intelligence? How do computers stack up against people? Computer scientists have actively investigated game playing for decades, as it provides a backdrop for understanding the learning process and ultimately determining if a computer system can possibly learn. Computer scientists have constructed systems that do learn over time. There are chess-playing systems that learn to play better chess with experience. A system that initially moves its queen to unprotected squares may lose its queen to the opponent. If the system loses the game, an analysis may reveal that losing the queen was a turning point in the game. The system will respond by remembering not to position its queen on an unprotected square; that is, it will have learned how to play better chess. Learning strategies are important to systems that need to exhibit intelligent behavior.

Is this learning? On the one hand, computer systems are certainly able to autonomously learn within a specific domain, such as the confines of a chess game. On the other hand, people can learn about things for which they have had no prior "programming." Consider the following scenario: A person knows how to play chess but is unfamiliar with the game of checkers. Checkers and chess share some similarities and some differences. A person knowledgeable in chess could certainly watch a few games of checkers being played and learn how to play the game. A chess-playing computer system could not merely "watch" a game of checkers and then play checkers; it would have to be re-programmed to play checkers. The goal, then, is to develop computer systems that can learn autonomously, without outside intervention or re-programming.

In 1997, a chess-playing computer system, IBM's "Deep Blue," beat the reigning world chess champion, Garry Kasparov. IBM's supercomputer relied partly on AI techniques and partly on brute-force computational speed to play world champion chess. The central chess-playing algorithm in Deep Blue is an evaluation function that assigns a numerical ranking to each possible move and resulting board position. At first, it may appear that chess should be a simple game for a computer. The computer merely lists all possible moves and then chooses the best one. Unfortunately, there are too many possible moves to compute in a reasonable amount of time. Deep Blue employed a combination of expert intelligence (known moves, piece values, valuable regions, etc.) along with the ability to analyze 200 million moves per second. Because there are too many possible paths to search, Deep Blue employs a selective, rather than a brute force, search function. "Promising" paths are identified and followed while "unlikely" paths are ignored. The "intelligence" is found in the selection function, which is partly pre-coded and partly a learned response.

Were the designers of Deep Blue world champion chess players themselves? The answer is no. Some believe that the creators of an exceptional computer chess-playing system must have been exceptional chess players, but this is fallacious. The computer scientists designing Deep Blue certainly understood the game of chess. However, the system they created played better chess than the designers because of the hardware capabilities (raw speed), coupled with the software abilities (intelligent search algorithms). Again, the productivity and efficiency advantages of computerized problem-solving are clearly demonstrated. These advantages were the direct result of the human intelligence behind the problem solving process.

Fifteen years after Deep Blue, another IBM supercomputer, Watson, beat the best human beings at the game of *Jeopardy!*, a game show that requires a deep understanding of language. The object is to correctly state the question related to the answer that is given. According to IBM, Watson is a cognitive system that understands natural language.

> "*Jeopardy!* was selected as the ultimate test of the machine's capabilities because it relied on many human cognitive abilities traditionally seen beyond the capability of computers, such as:
>
> - the ability to discern double meanings of words, puns, rhymes, and inferred hints;
> - the capacity for extremely rapid responses;
> - the ability to process vast amounts of information to make complex and subtle logical connections.
>
> In a person, these capabilities come from a lifetime of participation in human interaction and decision-making, along with an immersion in pop culture."[3]

Intelligent Computers

Will future computers be intelligent? Yes. Admittedly I can't be certain of my answer, but I do believe that computers can be intelligent. One reason for my optimism is the creative and innovative spirit of human beings.

Bill Gates in his book, *The Road Ahead*, relates a story that is probably apocryphal, but nonetheless, enlightening. According to this story (which others have identified as a myth), the head of United States Patent Office declared in 1899 that the office should close because everything that could possibly be invented had already been invented.[4] Of course, there have been a few new inventions since the end of the nineteenth century! Human beings are wonderfully creative, and I'm not willing to bet against human ingenuity.

Some Christians are startled when I claim that computers can possibly be intelligent. There is certainly nothing unbiblical about the possibility. God created intelligent entities, including human beings. As the pinnacle of His creation, we have been endowed with a (tarnished) image of Him. Part of that image is reflected in our creativity. Computer systems already *appear* intelligent. The possibility of human beings creating truly intelligent systems seems likely. Notice, however, the vast difference between "intelligent computer" and "human equivalent system" (the objective of Strong AI research). An intelligent computer system would not be human equivalent. People are intelligent, but what makes us human is much more than mere intelligence; it is our soul and spirit created in the image of God that makes us human.

An intelligent computer system would not be human equivalent. People are intelligent, but what makes us human is much more than mere intelligence; it is our soul and spirit created in the image of God that makes us human.

Human Equivalence

Will you live on in the mind of a computer? This was the provocative title of one of the earliest popular-level accounts of Strong AI.[5] Will computers eventually be human equivalent? The key word in this question is "human." The answer to the question depends upon a worldview and how that worldview defines "people." This objective of strong AI research is vastly different from the "intelligence" objective. Because it involves worldview issues, there are a variety of opinions on the issue.

Computer scientist Hans Moravec, who Michael Hirsh quoted in his AP wire story, believes that computers will become human equivalent. Moravec believes that human consciousness is the result of naturalistic brain processes. The bundle of neurons known as the brain produces the mind and human consciousness. Essentially, the hardware of the brain is analogous to computer hardware, as both are built out of switches. According to Moravec, as soon as computer technology is advanced enough, it will be possible to capture human consciousness and download a person's living essence into a computer. Moravec's worldview profoundly influences his answer to the "human equivalence" question. Moravec's naturalistic, materialistic worldview is evident in his statement regarding origins: "We owe our existence to organic evolution. But we owe it little loyalty." According to Moravec, human beings are just material beings, and the brain is the result of a continuous, naturalistic process (evolution). If this were true, then it would indeed be possible to produce human equivalence in a machine.

Mathematician Roger Penrose disagrees with Moravec. Penrose's work has challenged the typical Strong AI view that the mind is produced from an interconnected system of neural networks. Penrose claims that Strong AI via present computer systems cannot in principle duplicate the workings of the human brain. Penrose argues that consciousness exists outside the realm of computability, as the human mind can conceive some problems which are not computable. There do exist entities (both in mathematics and "reality") which we know to be true but which cannot be proven or calculated. Penrose does not believe that an algorithmically-based computer can capture the human essence. Penrose's worldview is very similar to Moravec's; however, it is naturalism with a twist. For Penrose, quantum mechanics is the "missing link" that makes naturalism work. While Penrose is correct in stating that human consciousness is outside the realm of computability, he does not see that the "missing link" is something outside of the physical universe, namely a transcendent God. It is not the mysticism of "quantum mechanics" that accounts for the human mind; rather, it is the inherent "image of God" that is responsible for our inimitable minds.

Computer scientist Fred Brooks presents an alternative viewpoint. According to Brooks, the Maker (the Triune God) gave humanity a number of inherent attributes, one of which is the gift of sub-creation. Brooks refers to these unique human attributes as "birth day gifts" since they were imparted to the original human beings at their birth, as recorded in Genesis 1.[6] Sub-creation is the ability and call of creative, rewarding work. J. R. R. Tolkien casts light on the concept of human creativity as sub-creation in the poem "Tree and Leaf":

> Although now long estranged,
> Man is not wholly lost nor wholly changed,
> Disgraced he may be, yet is not dethroned,
> and keeps the rags of lordship once he owned:
> Man, Subcreator, the refracted Light
> through whom is splintered from a single White
> to many hues, and endlessly combined
> in living shapes that move from mind to mind.
> Though all the crannies of the world we filled
> with Elves and Goblins, though we dared to build
> Gods and their houses out of dark and light,
> and sowed the seed of dragons—'twas our right
> (used or misused). That right has not decayed:
> we make still by the law in which we're made.[7]

As Tolkien reminds us, we are creative because we were fashioned by a creative God, and we still maintain an image of our Creator, although now tarnished as the result of sin. While Tolkien demonstrates the gift of sub-creation in the creation of fantasy, Dorothy Sayers extends the gift of sub-creation to all forms of human creativity. Sayers' work, *The Mind of the Maker*, demonstrates a "division of labor" for creation.[8] Sayers demonstrates the role of each Person in the Trinity in creation: Father as Idea; Son as Energy; and Spirit as Power. The Father conceived the creation Idea, envisioning the whole from beginning to end, even before it physically existed. The Son provided the Energy to call the creation into existence. The Spirit's Power enables us to interact with and understand the creation. This three-fold concept of Idea, Energy, and Power is reflected in how people sub-create in literature and the arts.

Brooks ultimately extends the gift of sub-creation to the work of computer science. Rather than invest time and energy in creating AI, Brooks argues for an investment in IA (Intelligence Amplifying). Brooks believes that Strong AI sent the discipline of computer science off in a wrong direction. Brooks' thesis is that IA > AI; that is, an Intelligence Amplifying (IA) system can better any AI system. For example, even though Deep Blue beat Garry Kasparov in chess, if we arm Garry Kasparov with a sophisticated IA chess playing system, the combination of the human and the computer (IA) will surely beat the computer (AI) alone. Brooks identifies computer scientists as toolsmiths and claims that their delight is found in fashioning power tools and amplifiers for users' minds. Brooks's Christian worldview leads him to the proper conclusion, and he rightfully decries the tremendous waste of money and human talent in the pursuit of human equivalence.

Because the "human equivalence" goal of strong AI depends upon a worldview, we must find the *true* worldview as revealed in the Bible in order to lay the foundation to answer the question. What is a human being, and who defines what people are? There are radically different answers depending upon the worldview framework employed. According to a humanist worldview, people are merely physical entities whose existence is a cosmic accident resulting from a naturalistic, evolutionary origin. If this worldview were true, then the strong AI goal would be achievable. If people are just a bunch of organic switches (the brain is just a computer made out of meat), then it is certainly possible to capture the essence of human beings in a computer system. According to the Christian worldview, people were specially created in the image of God. People are not just physical entities but possess a soul and spirit that reflect that now tarnished image.

The goal of strong AI research is eternal life. Consider these quotes from the AP wire story[9]:

> "If you can survive beyond the next 50 years or so, you may not have to die at all—at least, not entirely."

> "In an astonishingly short amount of time, scientists will be able to transfer the contents of a person's mind into a powerful computer, and in the process, make him—or at least his living essence—virtually immortal, Moravec claims."
>
> "MIT's Gerald J. Sussman, who wrote the authoritative textbook on artificial intelligence, agreed that computerized immortality for people "isn't very long from now."

Will you live on in the mind of a computer? The answer is no. God is the author and creator of life. Only God can create a human soul. The sad part of the story for those who accept a naturalistic worldview and have placed their hope in Strong AI is that eternal life is already available! As Jesus declares: "For God so loved the world that he gave his one and only Son, that whoever believes in him shall not perish but have **eternal life**" (Jn 3:16, emphasis added).

So, we have a great message to share. Our hope is not in AI but in the Creator.

> Will you live on in the mind of a computer? The answer is no. God is the author and creator of life. Only God can create a human soul.

Conclusion

Computers can be intelligent. While many people believe computers are already intelligent because they appear intelligent, the possibility of computers being truly intelligent is real.

Even if computers are intelligent, it does not mean they can be human equivalent, since intelligence does not equate to "human being." There is much more to being human than merely intelligence. At the forefront is the fact that people were created in the image of God.

Christians can use computers as powerful problem-solving tools. If, or when, computers become truly intelligent, they will be even better tools. The Christian mission of making disciples is more important than ever. While some will falsely cling to the hope of a man-made eternal life, we have a timeless message that can bring true hope to the world. Thank God that eternal life does not depend upon a bug-ridden AI system created by fallible human beings! Jesus declared, "I am the way and the truth and the **life**. No one comes to the Father except through me" (Jn 14:6, emphasis added).

Endnotes

[1] See http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html and http://www.livinginternet.com/i/ii_ai.htm.

[2] Edmund Callis Berkeley, *Giant Brains; or, Machines that Think*, 6th ed. (New York: Wiley, 1949), 5.
[3] http://m.ibm.com/http/www-03.ibm.com/innovation/us/watson/science-behind_watson.shtml
[4] William Gates, *The Road Ahead* (New York: Viking, 1995).
[5] Michael Hirsh, "Will You Live on in the Mind of a Computer?", AP wire story, June 1987; original source: June 14, 1987 edition of *Milwaukee Journal*, page 1J.
[6] Fred Brooks, "Computer Scientist as Toolsmith II", *Communications Of The ACM* 39, no. 3 (March 1996), http://www.cs.unc.edu/~brooks/Toolsmith-CACM.pdf.
[7] J. R. R. Tolkien, *On Fairy Stories* (New York: HarperCollins, 2008).
[8] Dorothy Sayers, *The Mind of the Maker* (San Francisco: HarperCollins,1987), 35–45.
[9] Hirsh, "Will You Live on in the Mind of a Computer?"