

FROM THE ARCHIVE / 1990

Designing the Human Mind

Robin Martherus | BYU-Hawaii | 1990

This is a photographic reproduction of an undergraduate paper rediscovered in 2026. It preserves the original text, illustrations, notes, and bibliography, including limitations and conclusions the author would qualify today.



The original cover used a student ID instead of the author's name so reviewers could evaluate the paper without knowing the author and potentially introducing bias. The ID has been masked in this public copy. This is the only intentional alteration to the photographed content.

The photographs are presented in reading order: original cover, abstract, pages 1-13, notes, and two pages of works cited. One duplicate photograph of page 1 has been omitted. The original photographs remain unchanged.

The bibliography records a BYU-Hawaii student survey conducted March 5, 1990, and a personal interview on March 3, 1990. The paper reflects the research the author was studying at the time; it is not presented as a prediction of modern AI.

Companion article: [Thirty-Six Years Later, the Question Has Become Operational](#).

Designing the Human Mind



Abstract

The ultimate goal of artificial intelligence is the ability to duplicate the human mind. Several things need to be taken into consideration in order to realistically attempt this great task. Among these things are the definition of intelligence. Using the dictionary definition of intelligence and several simple tests we can determine that computers can indeed be intelligent. Some differences between human minds and computers are the mind's ability and the computer's disability of parallel processing and the way that humans and computers interpret symbols. Among the attempts to close the gap between computers and human minds come neural networks. These computer programs are actually capable of generating brain-waves. As we continue to study artificial intelligence in its purest form, it becomes a subject of philosophy rather than technology. Because artificial intelligence is a subject of philosophy, we cannot come to a conclusion as to whether it is actually possible to duplicate the human mind.

Designing the Human Mind

Computers have been a part of our lives for several years. Every day computers seem to be able to do more than the day before. Each day they are able to perform an increasing number of functions that previously could only be performed by humans. That brings up many interesting questions: Will computers eventually be able to perform the same complex tasks that our human mind can do? Is it possible that computers could be programmed to make decisions based on a conscience or a set of values? The answers to these questions are not so easily realized.

Modern medicine is not sure how humans are able to make decisions. Therefore, modern men can not be sure whether or not we can duplicate these procedures without first understanding the natural processes in the human brain. We do have some information available regarding the thought processes of our brain. Taking these individual processes, analyzing them, and eventually duplicating them using a computer can help us better evaluate the feasibility of the eventual duplication of a complete set of artificial, human emotions.

We will begin by analyzing the concept of intelligence with regard to humans and machines. Computer scientist Marvin Minsky made the following prediction about machines and their ability to display intelligence in Life magazine 19 years ago:

In from three to eight years, we will have a machine with a general intelligence of an average human being . . . a machine that will be able to read Shakespeare, grease a car, play office politics, tell a joke, have a fight. At that point, the machine will begin to educate itself with fantastic speed. In a few months, it will be at the genius level, and, a few months after that, its powers will be incalculable. (Carpenter 64)

Of course that statement was made a long time ago, and most of us laugh at it because no such thing has happened. But, here is a similar statement made only a few years ago. Professor Edward Fredkin, the manager of the Artificial Intelligence Laboratory at MIT, stated the following during an interview broadcast on BBC television in October 1983:

There is no principle of science or engineering that prevents us from making intelligent computers that are infinitely smarter than ourselves. . . . We must ask ourselves what such machines will be doing in the future. . . . They may take away some of our (nuclear) toys, they will solve weighty problems that we ourselves have been unable to solve. They will talk to us only to amuse themselves and so, in some cases, keep us as pets. (Aleksander and Burnett 194)

Fredkin does not say that computers will be able to do these fantastic things within a few years, but does make a bold

statement regarding the possibility of such a breakthrough in the future. Obviously, computers have not yet taken over our lives to the extent that they have made us their pets (at least not to our limited knowledge). But, computers can, in some ways, be considered infinitely smarter than ourselves. The word "smarter" implies greater intelligence.

Because one of the main aspects of human thinking is intelligence, we must discuss the definition of intelligence and decide if computers are truly capable of being intelligent. The Random House Dictionary defines intelligence as the "capacity for reasoning, understanding, and for similar forms of mental activity." It also defines intelligent as "having, displaying, or characterized by intelligence or a high mental capacity" (460). Several ways have been developed to determine whether machines are able to show intelligence or obtain a high mental capacity.

Igor Aleksander and Piers Burnett state "The ability to solve problems of one sort or another is widely used as a yardstick of intelligence . . ." (Aleksander and Burnett 58). They go on to say that all a researcher has to do is choose a task that seems to require intelligence (such as playing chess or recognizing images) and try to build a machine that could accomplish this task. If the machine is successful then it will have shown itself to be intelligent (13).

Alan Turing, a British mathematician, devised a simple test to determine whether a machine is intelligent or not. George Johnson, author of several books in the field of

artificial intelligence, explains that this test is composed of a judge communicating by Teletype¹ with a human and the machine to be tested. The judge poses a series of questions to both the human and the machine. If, from the answers, the judge cannot tell which is the human and which is the machine then the machine is considered to be intelligent. (To pass the test, Turing wrote, the machine might have to occasionally give wrong answers and pretend to take a little longer than necessary to perform difficult calculations to seem more human) (Johnson 44).

Of course, most computers of today would be considered intelligent using such a test. We must therefore look a little closer to the definition of intelligence. Although a computer passes the Turing test for intelligence, does the computer understand the questions? We might argue that if the computer gives the right answers it automatically has to understand the question. The other half of the definition for intelligence is that the computer must be able to reason.

The Random House Dictionary defines reasoning as "The mental powers concerned with forming conclusions or inferences" (731). Again it is obvious that the computer must have the capability to form conclusions or inferences by the simple fact that it came up with the right answers. The process that allows a computer to make conclusions is defined by a computer program.

A computer program is a set of instructions that the computer must follow to reach an ultimate answer to a

problem. The program contains all the rules and directions that must be followed by a computer in order to acquire the desired results (in some cases the right answer to a question.)² These programs are written in special computer programming languages that both the programmer and the computer can understand.³ The computer program can make the computer perform tasks that require intelligence, therefore, making the computer intelligent. Since both humans and computers exhibit intelligence, it cannot be used to differentiate between the human mind and a computer. There must be something else that does.

James Sacra Albus, a leading researcher of artificial intelligence, stated the following regarding differences in problem solving:

Perhaps the major difference between human and machine approaches to problem solving is that the brain performs many parallel operations while computers do not easily represent parallel processes in search strategies. (Albus, James S. 284)

Parallel processing means the ability to do more than one thing at the same time. At the present time there is no computer that is able to function in this way. There are some computers, aptly named "super computers", that attempt parallel processing by having multiple "brains" or central processing units (CPUs).⁴ One example of such a computer is the Cray2.

Dr. Gordon Hoagland, professor of computer science at Ricks College in Rexburg, Idaho explained during a lecture in an advanced computer class how the Cray2 works. The Cray2 computer simulates parallel processing by assigning a different process to each CPU. Each individual CPU runs its assigned process at the same time that the other CPUs run theirs, therefore, creating parallel processing (Hoagland). In order to simulate all of the processes running simultaneously in the human brain, we would have to have an enormous number of CPUs. We do not have the technology to sustain the number of CPUs that would be required to simulate the human brain. Another type of computer that simulates parallel processing is a multi-user computer.

Many multi-user computers are in use today. Some common ones are Digital's Vax, Honeywell's Ultimate, and the Prime Computer. Dr. Ernest Carey, professor of computer science at Brigham Young University-Hawaii Campus, explains multi-user computers as follows. He states that these computers allow a number of people to use the same CPU by the use of a system called time-sharing.

A time-sharing CPU won't run several processes at the same time, but will complete a small portion of one process before completing a small portion of another process and then returning to the first process to complete another small portion and so forth. Because computers are very fast, a time-sharing system gives the illusion of parallel processing (Carey). In a time-sharing system, the more processes the

CPU is working on, the longer it takes for one process to be completed.⁵ Today's technology does not allow a computer to be fast enough to be able to handle the number of processes that the brain is able to complete simultaneously in the relatively short time that it takes. We have now found one of the obstacles that prevents us from reaching our goal of duplicating the processes. Unfortunately, there are several more obstacles to overcome that will require some rethinking in the way we program computers today. Let us return to the assumption made earlier, during the Turing test, that the computer must have understood the question in order to have given a right answer.

There is another test during which both humans and computers will score the same, but will illustrate a critical difference between the computer's way of understanding and the human's way of understanding. This test is called the Chinese Room Paradox. John Searle explains the test this way:

... suppose I am placed in a room containing baskets full of Chinese symbols. Suppose also that I am given a rule book in English for matching Chinese symbols with other Chinese symbols. The rules identify the symbols entirely by their shapes and do not require that I understand any of them. ... Imagine that people outside the room who understand Chinese hand in small bunches of symbols and that in response I manipulate the

symbols according to the rule book and hand back more small bunches of symbols. (Searle 26)

Searle goes on to explain that if the bunches that are given him are "questions," and the bunches he returns are "answers," and the answers are indistinguishable from answers given by someone who understands Chinese, then Searle would pass the Turing test for intelligence but lacks the understanding required by the dictionary's definition of intelligence. The difference in Searle's level of understanding and the level of understanding of those who understand Chinese is the concept of syntactic versus semantic.

Searle describes the difference as follows. Computers manipulate symbols without meaning. These symbols are purely formal (syntactic). Computers follow a strict set of rules for handling these symbols (programs). The human mind actually attaches meaning to each of the symbols before manipulating them [see illustration 1](Searle 27). "What makes this difference so important is that symbols are purely abstract notions: they have no essential physical properties to define them and can be implemented in any physical medium whatsoever"(Searle 27). Let's use this paper to clarify this essential difference between man and computer. The words on this paper were typed into a computer before printing. The author used symbols (letters and words) to relay an abstract idea to the reader. When we read this paper we can visualize the meaning of the words.

On the other hand, the computer only saw the words as a set of symbols. The computer has no idea what each of the words really mean.




Illustration 1 taken from Scientific American Jan '90 p28.

Consequently, we have found another obstacle in our search for artificial minds. Recent technology has tried to hurdle this obstacle by creating what is called artificial neural networks. Several artificial intelligence research

groups decided that because the center of human thought is located in the brain, they should first duplicate the brain before they can hope to duplicate the thought process. Keeping this in mind, they started to take apart the components of the brain in order to evaluate their function and analyze the possibility of duplication. What they came up with is what is called an artificial neural network. They took the science of neurons, dendrites, and other brain parts and duplicated them electronically (Arbib 19).

While trying to duplicate some of the processes within the brain, several other researchers ran across some interesting findings. Otis Port, an editor of Businessweek magazine describes the finding this way in a recent article:

It sounds like something out of a science-fiction film: A computer model of a portion of the human brain has become so realistic that it shows signs of "life." After recent improvements, the model began to generate its own brain waves -- spontaneously, without being told to do so. That stunned Dr. Roger Traub, the IBM scientist who began creating the model 10 years ago in collaboration with Columbia University researchers. They double-checked their work hundreds of times to be sure that the software does, in fact, exhibit unprogrammed behavior. (Port 84)

This finding opens up more possibilities that are yet to be

discovered. Not only can this research help us expand the horizons of artificial intelligence research, but it can help medical technology by finding the cause of these unexpected results. (the above article goes on to explain that the brain waves produced by the computer are the same as those found to trigger epileptic seizures in humans.) This research has shown that computers have the ability to exhibit behavior that previously was thought to be unique to living organisms. As a matter of fact, the medical profession considers a human to be alive while his brain is still producing brain waves. When the brain stops producing brain waves, a doctor can pronounce the person to be brain dead. According to this line of thinking, Traub's computer was alive! Although this finding seems to have no connection to the human thought process, it does show hints of things yet unknown. The "unknown" crosses the line between the domain of technology and the realm of philosophy.

Philosophy is the "study of the truths and principles of being, knowledge, or conduct." (Random House 660). We need to be concerned with the philosophical aspects of artificial intelligence, namely, what is the human mind? Is it an abstract idea or a physical object? If it is an abstract idea, then we must rearrange our computer programs to be able to handle abstract ideas instead of symbols that represent them. When dealing with abstracts we cannot be sure of anything and can only theorize. If the human mind is indeed a physical object, there should be no reason why future

technology will not find a way to duplicate it using computers. At this time, it seems as though we need to focus on the question of whether the human mind is a tangible object or not. If we can answer this question then we can continue our quest for the artificial mind.

It might be interesting to find out how the general public would answer the question of whether they thought that computers would be able to think like humans in the future. When asked, the majority of students surveyed at Brigham Young University - Hawaii Campus did not believe that computers would someday be able to think and reason like humans can (survey). Most seemed to have philosophical reasons to back their answers. John Searle answers the question this way: "No one expects to get wet in a pool filled with Ping-Pong-ball models of water molecules. So why would anyone think a computer model of thought processes would actually think?"(Searle 31). Another way of looking at it would be to assume that a computer simulation of the digestive system would actually be able to digest a slice of pizza. The computer could show how a digestive system would accomplish this task by representing the different ingredients of the pizza with symbols, but it would never be able to digest an actual pizza.

Paul and Patricia Churchland, professors of philosophy at the University of California at San Diego do not think that we need to duplicate the brain exactly in order to duplicate the human mind. They put it this way: "Requiring perfect

parity [in duplicating the human brain] would be like requiring that an artificial flying device lay eggs" (Churchland 37). In other words, we are not sure what parts of the brain are required for thinking. So, why duplicate the entire brain in order to generate artificial thinking? It has become obvious that the more research we do on artificial intelligence, the more questions are raised.

We have discussed many aspects of artificial intelligence that are essential in the eventual development of a true artificial, human mind. What then is the answer to the question of whether it is possible to create an artificial human mind? We have moved from a definite possibility to a definite maybe. At this time we can honestly say that we do not know. Only time will give the true answer to this controversial question.

Notes

1 Teletype is a brand of teletypewriter which is a typewriter similar to a computer keyboard that sends characters over a wire to another teletypewriter or a computer where it can be printed or analyzed.

2 The author of this paper is a seasoned computer programmer and can be considered an authority on computer programs.

3 Many programming languages exist and most are designed for specific applications. For example, Fortran (stands for FORMula TRANslation) was designed to aid in programming mathematical formulas. LISP was specifically designed for artificial intelligence.

4 A central processing unit or CPU is the main information processor of a computer. This is where all of the instructions and data are analyzed and processed.

5 The author of this paper has several years of experience using multi-user computer systems.

Works Cited

- Albus, James Sacra. Brains. Behavior. & Robotics.
Petersborough, NH: Byte Publications, 1981.
- Aleksander, Igor and Burnett, Piers. Thinking Machines: The Search for Artificial Intelligence. New York: Knopf, 1987.
- Arbib, Michael A.. Brains. Machines. and Mathematics. 2nd ed.
New York: Springer-Verlag, c1987.
- Brigham Young University - Hawaii Campus student body. survey conducted on 5 mar 1990.
- Carey, Ernest. Professor, Computer Science Division, Brigham Young University-Hawaii Campus. Personal interview. 3 Mar 1990.
- Carpenter, Betsy. "Will Machines Ever Think?" U.S. News & World Report Oct 17, 1988: 64-65.
- Churchland Paul M. and Churchland Patricia Smith. "Could a Machine Think?" Scientific American Jan. 1990: 32-37.
- Crawford, Michael. "drawing of Chinese Room." Scientific American Jan.1990: 27.
- Hoagland, Gordon PHD. Computer Science professor. "CS131." lecture. Ricks College. Fall 1984.
- Johnson, George. Machinery of the Mind: Inside the New Science of Artificial Intelligence. New York: Times Books, c1986.
- Port, Otis. "Uh, Doctor, I Think You Should Take A Look At This Computer." Businessweek 5 Mar. 1990: 85.
- Pylyshyn, Zenon. "Complexity and the Study of Artificial Intelligence." Haugeland, John. Editor. Mind Design:

Philosophy. Psychology. Artificial Intelligence. Cambridge, Mass: MIT Press, c1981: 67-94.

Random House Dictionary. The. New York: Ballantine Books, (1980)

Searle, John R. "Is the Brain's Mind a Computer Program?"
Scientific American Jan. 1990: 26-31.