Spring 1989

In the Image of Man: Reflections on Artificial Intelligence

Noreen L. Herzfeld
College of Saint Benedict/Saint John's University, nherzfeld@csbsju.edu

Recommended Citation
Herzfeld, Noreen. 1989. In the Image of Man: Reflections on Artificial Intelligence. *Symposium: CSB-SJU Faculty Journal* 7: 13-21.

In the Image of Man: Reflections on Artificial Intelligence

Noreen Herzfeld

"Then God said: 'Let us make man in our image, after our likeness. Let them have dominion over the fish of the sea, the birds of the air, and the cattle, and over all the wild animals and all the creatures that crawl on the ground.' God created man in his image; in the divine image he created him; male and female he created them." (Genesis 1:26-27)

We have viewed ourselves as being created in the image of God. And this has been important in defining our image of ourselves. In relation to God, we are the creature, he the creator. However, we also view ourselves as creators; some have interpreted the 'image of God' spoken of in Genesis 1 specifically to refer to the ability we share with our creator to create.

And we have created. We have created art, literature, scientific theories, whole civilizations. Much of what we create is in our own image. We see images of ourselves in our art, be it literary or visual. We create verbal images of ourselves in our philosophies and theologies. Yet what we create is generally in our image only partially, often only superficially.

Man as creator is explored in the greek myth of Pygmalian, found in Ovid's *Metamorphoses*. Briefly, Pygmalian was a young sculptor, handsome and sought after by the local maidens, but unenticed by any of them. According to Ovid, Pygmalian was shocked at the vices of human women or, at least, at their imperfections. So he set about sculpting, creating the "perfect woman" out of ivory. When he finished his magnum opus, he had the unfortunate fate of having been so successful in his work that he found himself in love with his creation. He lavished attention and affection on the unheeding statue. The story has a happy ending; the gods take pity on Pygmalian and breathe life into his statue and they live together happily ever after.

There are several interesting points about this story that I will be referring to later. For now, it serves to illustrate the deficiency in most of our creations. They lack life; they are like us only superficially. The character in a novel, the figure in a painting, are limited to certain words, to one pose. They have neither minds nor

[This paper was delivered as a Faculty Coloquium on April 19, 1989.]

wills of their own.

Enter the computer. The computer too is a creation in our image. And like our other creations, the image is obviously partial. Computers don't look much like us. But the computer is different in that the part of ourselves that it imitates is our mind. And we identify with our minds. This can be seen in the very fact that we call ourselves "homo sapiens." We are thinkers, and in our minds we find our identity. So here we have a creation of a different order. In thinking about the computer, particularly with respect to artificial intelligence, we seem to have created something that is truly like ourselves.

But are computers intelligent? "Artificial intelligence" is a commonly used term in our society, but many people probably have little idea what it really means. To many, it conjures up visions of R2-D2 or Hal (from the movie *2001*). I'd like to begin by briefly summarizing what artificial intelliegence is and where it seems to be going.

Artificial intelligence is a loosely defined collection of areas of research that attempt to mimic some narrow area of human endeavor. These major areas are natural language processing, pattern or visual recognition, game playing and expert systems. I will briefly explain what each of these entails.

Natural language processing and pattern recognition are attempts to develop for the computer the human attributes of aural and visual understanding; to allow computers to hear and see and understand what it is they have heard or seen. Successes in these areas have been limited. Computers perform in these areas by attempting to match sounds or patterns to a given norm or template. Flexibility is lacking. If we ask someone from Birmingham, Alabama and someone from Bangor, Maine to pronounce the same sentence, we may get wildly varying results, and the possibilities are close to endless. Now, we may not always be able to understand various accents or regional variations ourselves, but we have a much more flexible system for picking out recognizable elements and puzzling out the rest than anyone has yet been able to program into a computer. Leaving aside the difficulties of the spoken word, even our written language seems to have endless possibilities. So much of the meaning of words is derived from their context that it becomes virtually impossible to program a computer with all the possibilities. Consider the sentences "Time flies like an arrow." and "Fruit flies like an apple." One cannot even pick out what is the subject and what is the verb without understanding the context of the sentence. There have been attempts at writing automatic translating programs. The results have been more humorous than otherwise. For example, one program which attempted to take into account cultural biases in translating between English and Russian translated the Biblical passage "The spirit is willing but the flesh is weak" into "The vodka is good but the meat is rotten!"

The area of visual recognition suffers from similar difficulties. Consider for example, the task of recognizing a chair. There is no easy definition for chair that will take into account everything from straight-backed chairs, to beanbag chairs, to

Bauhaus design. A purely functional definition would have to include window ledges, car seats, and toilets. The task is difficult. We use an incredible amount of information simply to recognize a chair. Consider too the simple task of recognizing a face - something every six month old child is reasonably good at. The variety of expressions can make simple pattern matching a terrific chore for the computer. Features change in shape, size, and placement.

The computer has fared much better in the realms of game playing and expert systems. (For those unfamiliar with this term, an expert system is a program that attempts to answer questions or reason in an particular domain, as an expert does.) These are relatively easy areas for the computer primarily because they deal with extremely limited worlds. Game playing is a tailor-made situation for the computer because each game takes place in the limited world of the board and the moves allowed in the game, and these moves can be expressed using a finite list of rules. Chess, for example, is limited to a sixty-four square world with only thirty-two inhabitants, each with a very limited range of activities. These limitations allow the computer to grasp nearly all possibilities and so we have programs with master rankings. The real world, however, is infinitely more complex.

Expert systems involve similar limited worlds. There are expert systems, for example, to diagnose disease or analyze chemical bonds. (The business world has recently been pumping money into expert system research, hoping for one that will analyze Wall Street!) As an example of a successful expert system, I will describe MYCIN, a program developed by the medical school at Stanford to diagnose and recommend treatment for blood infections and meningitis. The patient's symptoms are entered into the computer and the computer can ask questions of the doctor to gain relevant lab or other information. MYCIN has been quite successful at diagnosis on partial information. On average, it prescribed fewer drugs than a comparative group of doctors, and successfully diagnosed 69% of its cases, as compared to an average of 62% for the doctors. (Luger 324) But again, one should note the extremely limited world that MYCIN is working in. This is not even the world of the general practitioner. The expert system approach does not work for general intelligence. The general information that informs our choices is too large a database to enter into any present day computer.

Thus, when compared to human intelligence, we see obvious limitations to computer intelligence. The source of these limitations lies in the difference between how the computer functions and how the brain functions. We don't know exactly how the brain functions. We do know that the human brain contains over 10 billion neurons, each capable of storing more data than a single bit or circuit in the computer. Still, supercomputers are beginning to approach this memory capacity. Furthermore, the brain's neurons are fairly slow when compared to computer circuits; the brain processes data in milliseconds as opposed to nanoseconds. So far, things look pretty good for the computer. What gives the brain its advantage is the fact that the neurons in the brain are connected to each other in an incredibly complicated network and the brain is theoretically capable of processing information in every neuron at the same time. The computer, on the other hand, processes only one piece of information at a time in a step-by-step fashion. Thus, though it can process any given piece of information faster than the brain, it is much slower at complex or multifaceted tasks. The faster processing speed of the computer shows up on problems that necessitate a step-by-step solution; your calculator can do arithmetic calculations faster than your brain. But problems that involve searching massive databases, making connections between several pieces of information, or leaps of intuition are done faster by the human brain.

Yet there has been progress. Computers have gotten faster and more efficient every year. But today we are facing a limit, the speed of light. Past gains in speed and efficiency have come primarily from making computer circuits smaller, allowing computers to have more circuits and, therefore, larger memories and reducing the distance the electrical current has to travel to turn circuits on and off within the computer, making processing of information faster. The only other factor that slows the computer down is the friction of the electrical current as it travels through a circuit. Assuming new advances in superconductivity, friction will become negligible. But we can only make circuits so small (they are already microscopic) and then we will be limited by the speed of electricity, the speed of light. Superconductivity may give us a fourth generation of slightly more powerful computers than we currently have. After that, the only way to make computers much more efficient is to mimic the parallel structure of the brain.

Computer scientists are working on models called neural networks to do just that. However, there are problems. Current models have been quite small (on a scale of 100-200 processors) and performance declines exponentially as the number of processors increases. Currently, no one has any idea how to handle billions of processors simultaneously. The number and nature of the connections quickly exceeds the conceptual grasp of our conscious minds. Moreover, we don't know enough about how the mind works to get a good model. The question arises whether the human brain can ever fully comprehend its own workings. Can we understand mind with mind? As mathematicians and philosophers know, self-referential problems are fraught with difficulties and paradoxes. Currently, we seem to be bound by the fact that our conscious minds work in a step-by-step fashion and we project this method onto the computer. (For a complete discussion on neural networks, see Folsom.)

We are left with the fact that we are unlikely to make any real strides in artificial intelligence unless we manage to change the structure of the computer itself, and find a way to master massively parallel structures. Until then, the world will simply be too large and complex a thing to be comprehended in a step-by-step manner.

It is generally agreed that we would have a truly intelligent machine if it could pass a test devised by Alan Turing. This test is a version of the old TV game "To Tell the Truth." A man would be placed in one room and a computer in another. One could question both of them and would be asked to tell which was the computer and which the man. True artificial intelligence would imply that one would not be able to tell the difference. We are a long way from producing a machine that would not fail this test very quickly. The debate has been raging for years as to whether we will ever produce such a machine. Rather than enter into this fray, I would ask a different question: not "Can we produce such a machine?" but "Why would we want to?" Why do we want to create artificial intelligence?

This question has not been addressed seriously, to my knowledge, in any of the literature. I have the feeling that most computer scientists would find the answer self-evident. "Why not?" But the question cannot be so easily dismissed. After all, we're not talking about an isolated Dr. Frankenstein working away in the basement. The search for artificial intelligence is a multi-million dollar effort involving both the business and university research communities, not to mention the Pentagon. It would seem that a scrutiny of our motives is not out of order.

Perhaps the most obvious answer to the question of "why?" is that we look to the computer for convenience or power. We have always designed machines to do our work for us. Why not one to do our thinking for us also? This reason seems very plausible on first glance. We want computers to analyze everything from Wall Street to battle situations. We would like robots that can see and hear and bring us the morning coffee without bumping into every chair in the room. We'd like to sit back while our mechanical servants do the work, and gain for us power or wealth.

But carry this to its logical conclusion and we end up with Hal, a computer that is quite ready and willing to do our thinking for us. This is the stuff of science fiction nightmares. Norbert Wiener, in his last collection of essays, pointed out the problem with our desire for a machine to do our thinking for us. He likens this desire to our perpetual fascination with sorcery. We want to believe in the possibility of a Golem or a Genie; we want a servant to give us power by some means other than our own hard work. And to many people, the computer seems to give us this magical power. Magic goes on inside that little white box and answers come out that we didn't have to think up ourselves. But the folk literature of the ages warns us of the follies of relying on magical power.

Wiener uses one of the best stories of this genre to make his point, the story of the Monkey's Paw. Briefly, a family is sitting around the fire on a windy rain-filled night listening to the stories of a captain recently returned from India. At one point the son gets up to leave for the night shift at a local factory. The discussion among the others turns to native magic and the captain shows them a withered old monkey's paw, saying that it was endowed by a Hindu holy man with the power of granting three wishes to three successive owners. He says he does not know the first two wishes of the first owner, but that the third one was for death. He himself was the second owner, but his experiences were too horrible to relate. With that, he casts the paw into the fire. The host, of course, grabs the paw out of the fire and wishes for 200. There is a knock at the door. A somber foreman from the factory is there to inform the family that their son has been killed in an accident at the factory, and, although the factory owners themselves were in no ways to blame, they would like to offer the family 200 in remuneration and condolence (Wiener 59).

The family continues to make wishes that only make matters worse, but we can stop here. The moral is clear. Computers are as literal as the monkey's paw and though we might like them to think for us, chances are that we may not get what we asked for or, rather, we may get *exactly* what we asked for. Perhaps we had better to do our thinking ourselves. Clearly, some of us (including folks in the Pentagon) still want artificial intelligence for reasons of power or convenience, regardless of the dangers. But, admitting the dangers, as well as the moral shallowness of this reason, there are others who still continue to have an interest in artificial intelligence. They must have other reasons.

A second reason we may want artificial intelligence is simple curiosity. This is the "what can be done will be done" answer. By itself, this answer too seems insufficient. No one is spending millions of dollars trying to develop a non-prickly porcupine, yet it might be do-able. Lots of things could be done but aren't bothered with. Of course, the curiosity about artificial intelligence may be stronger than that for non-prickly porcupines; after all, artificial intelligence might tell us something about ourselves, a topic we are perennially interested in. Unfortunately, this has not proven to be the case. Artificial intelligence research has told us relatively little about the way we think. Rather, only by borrowing from the psychologists, have we learned a little about how to design artificial intelligence. Indeed, there is a danger in interpolating in the other direction. We might miss possibilities in ourselves by limiting our perception of our thought processes to what a computer is capable of doing. This limiting already can be seen when we consider data that can be quantified as more pertinent, more "scientific", simply because it can be entered into a computer.

Having summarily dismissed power, convenience, curiosity, and self-understanding as insufficient motives by themselves to drive us in our search for artificial intelligence, what is left? Let us return to the Pygmalian myth for a clue to another possibility. Why did Pygmalian carve his perfect woman? Why a woman? Why indeed did God create humankind in the first place? Why this drive to create in one's own image? I think one answer lies in our intrinsic loneliness. As Henri Nouwen points out, "loneliness is one of the most universal human experiences, [and] our contemporary Western society has heightened the awareness of our loneliness to an unusual degree." (14) He also suggests that "too often we will do everything possible to avoid the confrontation with the experience of being alone, and sometimes we are able to create the most ingenious devices to prevent ourselves from being reminded of this condition." (16)

I'm not suggesting that we want thinking computers merely as companions to fill the moments when we find ourselves alone. The loneliness we seek to have alleviated by artificial intelligence is much deeper and more fundamental. It is the loneliness of being the only rational creatures, a loneliness of the species rather than of individuals.

Nobel prize winning biologist, Jacques Monod, touches on this loneliness in his final work on evolution. As man begins to accept the theory of evolution, he must "wake to his total solitude, his fundamental isolation" in an uncaring universe. (172) Science has reduced our dependence on the existence of other rational beings. We no longer need God and the angels to explain the working of the universe. But how lonely it is without them! Our necessity to create some other type of rational being seems to be directly proportional to the waning of our faith in the previous existence of other rational beings.

The realization of our isolation as rational creatures brings with it both tremendous responsibility and tremendous anxiety. As Monod puts it, our destiny, our duties are nowhere spelled out for us but are of our own making. This is a big responsibility. We want to hold the power that being rational creatures gives us, but we do not want to bear the responsibility of that power alone.

We want to share responsibility. On a small scale, this can already be seen quite clearly in the use of computers in warfare. Joseph Weizenbaum of MIT has pointed out that policy-makers have "abdicated their decision-making responsibility to a technology that they do not understand—though all the while maintaining the illusion that they, the policy makers, are formulating policy questions and answering them." At the height of the Vietnam war, Admiral Thomas Moorer, chairman of the Joint Chiefs of Staff, was quoted in the *New York Times* as stating that the generals were "slaves to these damned computers," and that he could not help but base his decisions on "what the computer says." And no human being is responsible for what the computer says. (Weizenbaum 239) So many people contribute to the design, construction, and especially the programming of a large computer system that no particular person is ultimately responsible for the results.

There is one final point in the myth of Pygmalian that is illustrative. Pygmalian created out of lonliness and then fell in love with his creation. Are we in danger of doing this? Might we fall in love with the computer? Actually, some already have. We call them "hackers." They are a phenomenon that can be observed in the computer center on most large (and several small) campuses. Like a typical lover, the hacker is nocturnal, living for the hours he can spend with his love. He lives only for interaction with the computer and generally has very few human relationships. (I use "he" deliberately—this seems to be an predominantly male phenomenon. I have yet to meet a female hacker in this sense of the term, though I have known several men that fit this description. For a complete description of the compulsive programmer, see Weizenbaum, chapter 4.)

The hacker can be dismissed as a pathological case. On a larger level, one could suggest that our whole society has a bit of a love affair with the computer. And in so far as love blinds us to the shortcomings of the beloved, this could be a danger.On the other hand, we all seem to be fairly glad to believe that our creator has fallen in love with us; I would hope that if we somehow did manage to create intelligent computers that we would treat them with compassion.

In conclusion, we seem to be driven to create an intelligent machine, in our own image. This drive comes, at least in part, from our inner loneliness and our wish to share responsibility for this planet and indeed, for ourselves. But there are inherent dangers in our present course. These dangers come, in part, from the initial premise that what counts in our image is our ability to reason or think. But the image of God reflected in humankind is not simply our ability to reason, nor our ability to create, but our ability to love. And though we may love our creation in some way, we have not even considered endowing computers with the ability to love. Pygmalian's statue only came alive when it gained the ability to return the love its creator lavished on it. Whether we are capable of creating a machine with this ability remains to be seen. In the meantime, we are still in the realm of the superficial.

Works Cited

Dreyfus, Hubert. *What Computers Can't Do: The Limits of Artificial Intelligence.* New York: Harper and Row, 1972.

Folsom, Tyler. "The Search for an Electronic Brain: An Introduction to Neural Networks." *The CPSR Newsletter* 7:1 (1989): 14-21.

Luger, George and William Stubblefield. *Artificial Intelligence and the Design of Expert Systems.* Redwood City: Benjamin Cummings, 1989.

McCorduck, Pamela. *Machines Who Think.* San Francisco: W.H. Freeman, 1979.

Monod, Jacques. *Chance and Necessity.* New York: Random House, 1971.

Nouwen, Henri. *Reaching Out: The Three Movements of the Spiritual Life.* Garden City: Doubleday and Co., 1975.

Ovid. *Metamorphoses.* Trans. Rolfe Humphries. Bloomington: Indiana University Press, 1957.

Weizenbaum, Joseph. *Computer Power and Human Reason.* New York: W.H. Freeman, 1976.

Wiener, Norbert. *God and Golem, Inc..* Cambridge: M.I.T. Press, 1964.