

## HISTORY OF ARTIFICIAL INTELLIGENCE

P. McCorduck (Session Chairman), Univ. of Pittsburgh! M. Minsky, MIT: O. Selfridge, Bolt Beranek and Newman? H. A. Simon, Carnegie-Mellon University

### The Early History

P. McCorduck

Though our discussion is entitled "The History of Artificial Intelligence", in fact we are focusing here on one brief but highly significant moment in that history, the moment when art metamorphosed itself into science, from wish and dream to something like reality. As you will learn from each of the discussants, this metamorphosis took place at several locations during the early to mid-1950s, and its catalyst was the recognition that the computer was the most promising medium yet in which to realize what had been a human dream since earliest times, the creation of man-made, rather than begotten intelligence.

Let me remind you of some of the first manifestations of that dream. A great many artificial intelligences, or automata, appear in Greek mythology, put together to be useful or to carry out some task that the gods themselves find burdensome.

Around 850 B.C., Homer tells us about poor old ugly Hephaestus, the god of fire and the divine smith, who, because he is crippled, has to fashion attendants to help him walk and assist him in his forge!

These are golden, and in appearance like living young women.  
There is intelligence in their hearts,  
and there is speech in them  
and strength, and from the immortal  
gods they have learned how to do  
things.

From the immortal gods they have learned how to do things. There's a phrase so fraught with implications it takes your breath away. For humans to behave like gods--because godlike it is to imbue the inanimate with animation--is hubris indeed, and no opponent of artificial intelligence has failed to express shock at the blasphemous behavior of humans who aspire to divinity.

These opponents themselves, persistent and articulate as nearly everyone in the

field must know from personal experience, represent a tradition as ancient as the urge to create artificial intelligences.

Slightly before Homer was codifying what were surely already ancient traditions, another set of codes was brought down from Mt. Sinai by an unwilling prophet named Moses. We know these codes as the Ten Commandments, and it's the second one which is germane here: "Thou shalt not make unto thee any graven image or any likeness of any thing that is in heaven above or that is in the earth beneath, or that is in the water under the earth; Thou shalt not bow down thyself to them nor serve them, for I the Lord thy God am a jealous God..."

Indeed. No matter that the Lawgiver promptly violates that very commandment in the instructions for building the Ark of the Covenant; the message is clear. If you dabble in that sort of thing, you violate the territory of gods, and we all know who rushes in where angels fear to tread.

I like to think of these two attitudes as the Hellenic and the Hebraic. The Hellenic is curious, enthusiastic (a word which itself means filled with the breath of the divine) and generally at ease with the idea of artificial intelligence. The Hebraic, on the contrary, holds that the idea of artificial intelligence is fraudulent, wicked, and even blasphemous.

This is an arbitrary distinction when it comes to the actual business of doing artificial intelligence. Past and present, there are devout Jews and Christians untroubled by the idea. For example, among the ardent Christians in the past who figure in this history is Ramon Lull, a 13th century Spanish mystic who renounced the dissolute ways of his youth and went off to convert the Muslims (Cohen, 1966). It isn't recorded that he had much effect on them, but they had a profound effect on him: they introduced him to an Arabic thinking machine called a zairja, and he rushed back to Christendom with the idea of constructing a thinking machine of his own called, more grandly, the Ars Magna. (This translates to The Great Work of Art; Lull was right at home in a field not known for its humility.) The aim of the Ars Magna was to bring reason to bear on all subjects, and thereby arrive at truth without the trouble of thinking. Be that as it may, Lull's scheme seems to me remarkable not for its grandiose claims, but because without hesitation it presupposed that human thought could be mechanized.

Other well-known Christians were said to own brazen heads, that is, automata they had made themselves which were not only proof of their wisdom in being able to construct such things, but which then went on as consultants to amplify the wisdom of their creators. My favorite story is about the brazen head Albertus Magnus was said to own. "A lovely woman who could speak," says one source, and she so offended Albertus's pupil, the young Thomas Aquinas, that he burned it upon the death of his teacher (von Boehn, undated). What on earth did she say? Alas, the story loses some of its pi-quanty with the fact that Albertus out-lived his celebrated pupil by some six years.

The story of Rabbi Loew of Prague and his creation Joseph Golem is so familiar that it's hardly worth repeating: I merely want to remind you that the legend exists, and is all the more charming for the fact that several of the scientists associated with cybernetics and artificial intelligence have family traditions that trace their genealogy back to the Rabbi.

In short, my division between the Hellenic, or positive, or progressive, or irresponsible attitude (depending on your inclination) and the Hebraic, or negative, or backward, or responsible attitude (again, depending on your inclination) is merely a convenient way of illustrating that the two attitudes have coexisted with equal duration and intensity, which show no sign of abating.

In imaginative literature, the Hebraic attitude seems usually to have prevailed. Dr. Frankenstein found out to his chagrin what creating an artificial intelligence will get you, though the real story is more complicated than that, as are the issues. Later writers have been mostly pessimistic for the future of the human race side by side with artificial intelligences, by definition smarter, faster, and immune to human frailties. I don't know whether I count as pessimistic or optimistic Asimov's final story in his robot series, which gives us a paternalistic intelligence doing things for our own good and even making us like that state of affairs (Asimov, 1950).

You may begin to suspect that until the early 1950s, all the media in which artificial intelligences appear belonged to the realm of make-believe legend, fantasy, novel, play. If you classify scientific speculation as fantasy, this is probably true, but if by scientific speculation we mean not only a dream to be pursued but a possible means by which it can be accomplished, then we are in

different territory. In that case, first prize goes to Charles Babbage and his colleague the Countess Lovelace.

In 1843, Lady Lovelace published a long and detailed description of Babbage's Analytical Engine, and contrary to the implications of her widely quoted remark that machines can do only what we tell them to do, she added that the question of whether such an engine could be said to think would have to remain open until they actually constructed one and tried it out (Morrison and Morrison, 1961).

In any event, Babbage and Lady Lovelace considered building a quick chess machine in order to finance the building of the larger Analytical Engine, and were only dissuaded when they discovered that Tom Thumb was what the public was willing to pay to see, and not an automatic chess machine.

Later on, in 1915, two chess machines which played the endgame were constructed by Leonardo Torres y Quevedo, a gifted Spanish inventor. While he declined to claim that his automata were actually thinking, he suggested that we'd better refine our definitions of that process, and that his automata could certainly do many things which were popularly classified as thinking (Randell, 1973).

But the most passionate champion of machine intelligence was a man of breathtaking intelligence himself, Alan Turing. An intelligent machine might only be implicit in his famous proposal of the Turing machine in 1937, but nobody was more eager than he to make those implications explicit. He endured a lot of condescending derision for his dream, but he continued to pursue it, though the archives give the impression that except for one year, he was never able to pursue it as more than a serious hobby.

At the same time that Turing was at work, there was on the opposite side of enemy lines--this was by now World War II--a young engineer who had built the world's first up-and-running digital computer, installed in his parents' Berlin parlor. His name was Konrad Zuse, and he too was fascinated by the notion of intelligent machines. The possibilities of his machine's intelligence were clear in his mind: by 1943 he was wondering whether it could play a master in chess, and by 1945 he had developed a programming language called the Plankalkul which, he felt certain, could be used not only for mathematical problem solving but also for programming artificial intelligence problems of many kinds, though he believed that real artificial intelligence

was one or two generations away. Isolated by Germany's defeat and post-war prohibitions against electronic development, it came as a great shock to him to discover the mid-50s work of some of the people here.

In other words, the intelligent machine was an idea whose time had come, and it was not only that the computer presented a medium with which such a dream could be realized. There was a constellation of events, most notably the shift from one dominant paradigm, the physicist's notion of energy, to a new paradigm, the cyberneticist's notion of information, and there were the continuing efforts to describe psychological and biological phenomena in mathematical terms.

Because of these convergences, a young assistant professor of mathematics at Dartmouth College named John McCarthy, who himself had been fascinated by these issues for quite a while, suggested to his friends that some real progress could be made if only all of the people at work on these problems--all ten of them--spent the summer of 1956 together, helping each other. These three friends, who were Marvin Minsky, another young scholar who was a Harvard Junior Fellow in mathematics and neurology, Nathaniel Rochester, manager of information research at IBM's research center in Poughkeepsie, N. Y., and Claude Shannon, then a mathematician at Bell Laboratories who had much indeed to do with the paradigm shift from energy to information, agreed that it might not be a bad idea, and joined McCarthy in submitting a proposal to the Rockefeller Foundation for "a two-month ten man study of artificial intelligence to be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that 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."

This was the first time the term artificial intelligence had been used officially. John McCarthy won't swear he hadn't heard it before, but he was the first to apply it to the kind of work which was going on in this field, he promoted the term, and despite some other proposals and certain grumbling, artificial intelligence has stuck.

Rockefeller provided \$7500 and the initial four invited others who shared their faith. Among them were Trenchard More, Arthur Samuel of IBM, Oliver Selfridge and Ray Solomonoff of MIT, and two vaguely known persons from the RAND Corpora-

tion and Carnegie-Tech in Pittsburgh named Allen Newell and Herbert Simon.

After all this time, no one is quite certain how the Cambridge people got in touch with the Carnegie-RAND group, though there are several possibilities: Oliver Selfridge had given a talk at RAND the previous fall, and had mightily impressed Allen Newell, indeed, had turned his scientific life around. Marvin Minsky was a consultant at RAND and might have known about the work of Newell and Simon that way.

In addition, others came for short visits to talk about related work. Among these visitors was Alex Bernstein, then a programmer for IBM in New York City, who was invited to talk about the chess playing program he was working on, a program which was to receive a lot of subsequent publicity, to the horror of IBM, which feared that the idea of intelligent machines would be so threatening it would keep customers from buying computers.

If I were to share some of the things the Dartmouth Conference was supposed to accomplish, you might be tempted to laugh. John McCarthy recently took a look at the old proposal and he laughed, and suggested that by changing a few names and dollar amounts, the proposal might well be submitted today, more than twenty years later, and get a serious reading. He jests. Nobody really expected to accomplish all the things on the agenda for that summer, but neither did anyone intend to map out his professional life for the next twenty years, which in some cases is what happened.

This must be a very parsimonious version of the early history of artificial intelligence, and I have hardly attended to those who held--and still hold--that artificial intelligence is impossible, undesirable, and not worth the energy spent on it. The opposition isn't composed entirely of cranks--among the skeptics was John von Neumann--and only time will tell who is right.

But it seems to me a fine thing that some of the greatest visionaries, geniuses and crackpots of the western world have put their hand to the task of man-made intelligence. We sometimes forget that most scientific fields began with ideas that seem a bit loony to us now, and as a field takes on respectability, it would prefer to forget its disreputable antecedents. If we detect lunacy among the earliest forerunners, we had better admit that it is our very own, and here to stay. It is all of us humans who harbor that mysterious but ancient

urge to reproduce ourselves in some essential but extraordinary way. Artificial intelligence comes blessed with one of the richest and most diverting histories in science because it addresses itself to something so profound and pervasive in the human spirit.

Meanwhile, my story has come to the place where you can have eyewitness reports. To aid us all in understanding, I've asked these distinguished eyewitnesses to address some of the following questions:

First, when did the paradigm shift happen for you? Put another way, what did you think you were doing before artificial intelligence? What did that field look like? Why wasn't it sufficient?

Second, what is it that made the computer the instrument of choice? Were you looking for it, or did it pop up just in time?

Third, what made one project or problem more attractive than another at the time? Why did certain kinds of projects seem more appropriate at certain times? What were the canons of taste?

Finally, what should have been the strategy for AI in the first few years after the Dartmouth Conference in 1956?

Transcripts of the panelists' remarks will be published later in an appropriate journal.

### References

1. Asimov, Isaac, L. Robot. New York: Gnome Press, 1950.
2. Cohen, John, Human Robots in Myth and Science. London: Allen & Unwin, Ltd., 1966.
3. Morrison, Philip and Emily Morrison, Charles Babbage and his Calculating Engines. New York: Dover Books, 1961.
4. Randell, Brian, The Origins of Digital Computers: Selected Papers. Berlin: Springer-Verlag, 1973.
5. von Boehn, Max, Puppets and Automata. New York: Dover Books, undated.