Photo courtesy of MIT News Office

Joseph Weizenbaum (1923–2008)

Joel Moses and Jeff Meldman,
Massachusetts Institute of Technology

Joseph Weizenbaum was born in Berlin in 1923 and died there on 5 March 2008 at the age of 85. His German-Jewish family left Germany in 1935 and came to the US. Joe made major contributions to computer science and to the applications of computers, but was best known as a critic of AI. While working for General Electric in the 1950s, he helped develop one of the first banking applications using the magnetically encoded fonts on checks.

In 1964 Joe became a visiting professor and later a professor in the Massachusetts Institute of Technology's Electrical Engineering Department. At different times, each of us became a graduate student under Joe's mentorship. Later, we both became his colleagues. Those were exciting times at MIT in computer science. Project MAC began operations in 1963. It was a major research program that led the development of Multi-Access Computers, one definition of MAC. The Artificial Intelligence Group also joined Project MAC, working on Machine-Aided Cognition, another definition of MAC. Many students felt that MAC stood for Man against Computer.

The 1960s was an exciting decade intellectually for Joe. He developed the SLIP (*Symmetric List Processing*) language, a competitor of Lisp. SLIP differed from Lisp in a number of ways. It had both forward and backward pointers—hence, the "Symmetric" in its name.

SLIP also differed from Lisp in how it performed garbage collection.

In terms of impact, Joe's major achievement in the '60s was the Eliza program. Eliza, which was written in SLIP, accepted scripts and then a set of inputs from a human user. It translated the inputs to SLIP lists. Each script consisted of a set of rules, which transformed the input. There were many such scripts, but the most famous was a script which led to a program called "doctor" that generated responses resembling those of a Rogerian psychiatrist.

Joe said he wrote most of the script for the doctor program during a plane ride. There were a few dozen rules. Many of those rules looked for a keyword, such as "mother." Once Eliza identified the keywords or phrases in a sentence, it transformed them, keeping the rest of the sentence, and typed out the transformed sentence. If Eliza couldn't identify a keyword or phrase, it used a stock phrase, such as "Please go on." The "doctor" had no other understanding of what the user was typing.

At the time, circa 1966, few programs of any kind allowed conversations with users in a natural language. The doctor program embedded in Eliza was an instant success in attracting people. Joe was surprised when he caught his secretary using his program and telling it confidential things about her boyfriend.

Each of us remembers visiting Joe on different days at the Stanford Center for Advanced Studies in the

Behavioral Sciences during the 1972–73 academic year. Joe was on sabbatical, writing *Computer Power and Human Reason* (W.H. Freeman & Co., 1976). He had become concerned about the earlier work of the Stanford psychiatrist Kenneth Colby. The doctor script intentionally imitated a Rogerian psychiatrist who attempts to get the patient to keep talking about issues, such as the patient's relation to his or her family. But this was just an effort to mimic a human psychiatrist's conversation. Colby was apparently using a variant of the doctor script to actually treat paranoid patients.

The book dealt with a broader set of concerns than the one Joe had with Colby. Joe was also concerned with the attitudes of some AI leaders. He called these people the "artificial intelligentsia." There's no question that the claims made in the '50s by some, such as Herb Simon, about what computers would do were inflated. What was difficult for many of our MIT AI colleagues to accept was Joe's implication about the morality of AI research. Surely, he was right that asking computers to become judges in place of human judges isn't wise and might be immoral. But the state of the art in AI, some felt, was likely decades from the point that such an idea could even be reasonably contemplated.

Yet what truly concerned Joe ran deeper than what computers could or couldn't do—or should or shouldn't be permitted to do. His fundamental concern was that the prevailing inflation and uncritical trumpeting of AI achievements had the dangerous potential to produce a corresponding trivialization and devaluation of the human mind, of human responsibility, and of human experience as a whole. And he was understandably troubled by what he saw as the small number of fellow AI researchers who seemed to share this concern.

Terry Winograd was an MIT AI student in the early '70s and became a computer science professor at Stanford. He presented Joe with the 1988 Norbert Wiener Award of the Computer Professionals for Social Responsibility (CPSR). Wiener, who founded the field of cybernetics, was arguably the most famous MIT mathematician ever and was author of the 1950 book *The Human Use of Human Beings* (Da Capo Press). We quote a paragraph from Terry's tribute to Joe:

> From the point of view of CPSR, Wiener may be the patron saint, but Weizenbaum had much more direct influence on the fact that we are here tonight. During his many years of working with students at MIT, he was a teacher to many of us, and his work stimulated the thinking of many others who were not fortunate enough to be in the same institution. I know that my own concerns with social issues and the ethics of computing were strongly influenced by my contacts with Joe, beginning 20 years ago. All of us can trace some part of our concern back to Joe's vital influence.

We wonder to what extent Weizenbaum's German background influenced his views. We know that Herb Simon was a logical positivist and relatively close to analytic philosophy, which is popular in English-speaking countries. Continental philosophy tends to be more interested in the essence of things and less in pure logic. Weizenbaum is asking, what is the essence of our humanity? That question cannot and should not be dismissed. ◼

Joel Moses is an Institute Professor and a professor of computer science and engineering systems at the Massachusetts Institute of Technology. Contact him at moses@mit.edu.

Jeff Meldman is an associate dean for undergraduate education at the Massachusetts Institute of Technology and senior lecturer and director of undergraduate programs at MIT's Sloan School of Management. Contact him at jmeldman@mit.edu.