

Martin Lam gives us a British civil servant's view of the Lighthill report and subsequent developments. My comments concern some limitations of this view that may be related to the background of the author—or maybe they're just a scientist's prejudices about officials.

Lam accepts Lighthill's eccentric partition of AI research into Advanced Automation, Computer-based Studies of the Central Nervous System and Bridges in between. This classification wasn't accepted then and didn't become accepted since, because it almost entirely omits the scientific basis of AI.

AI didn't develop as a branch of biology, based on either neurophysiological or psychological observation, experiment and theory. It also isn't primarily engineering, although an engineering offshoot has recently developed. Instead it has developed as a branch of applied mathematics and computer science. It has studied the problem of systems that solve problems and achieve goals in complex informatic situations, especially the *common sense informatic situation*. Its experiments and theories involve the identification of the intellectual mechanisms, the kinds of information and the kinds of reasoning required to achieve goals using the information and computing abilities available in the common sense world. Sometimes this study divides up neatly into heuristics and epistemology, and sometimes it doesn't. Even connectionism, originating in a neurophysiological metaphor, bases its learning schemes on mathematical considerations and not on physiological observation.

Lam's inattention, following Lighthill, to the scientific character, goals and accomplishments of AI goes with a narrow emphasis on short range engineering objectives. Maybe this is normal for British civil servants. Nor is Lighthill the only example of a physical scientist taking an excessively applied view of scientific areas with which he is unfamiliar and finds uncongenial.

The Lighthill Report argued that if the AI activities he classified as Bridge were any good they would have had more applied success by then. In the 1974 Royal Institution debate on AI, I attempted to counter by pointing out that hydrodynamic turbulence had been studied for 100 years without full understanding. I was completely floored when Lighthill replied that it was time to give up on turbulence. Lighthill's fellow hydrodynamicists didn't give up and have made considerable advances since then. I was disappointed when BBC left that exchange out of the telecast, since it might have calibrated Lighthill's criteria for giving up on a science.

My own opinion is that AI is a very difficult scientific study, and under-

standing intelligence well enough to reach human performance in all domains may take a long time—between 5 years and 500 years. There are fundamental conceptual problems yet to be identified and solved, so we can't say how long it will take.

Many of these problems involve the expression of common sense knowledge and reasoning in mathematical logic. Progress here has historically been slow. It was 150 years from Leibniz to Boole and another 40 years to Frege. Each advance seemed obvious once it had been made, but apparently we earthmen are not very good at understanding our own conscious mental processes.

An important scientific advance was made in the late 1970s and the 1980s. This was the formalization of nonmonotonic logical reasoning. See (Ginsberg 1987). Not mentioning it in discussing the last 20 years of AI is like not mentioning quarks in discussing the last 30 years of physics, perhaps on the grounds that one can build nuclear bombs and reactors in ignorance of quarks. Logic needs further improvements to handle common sense properly, but no one knows what they are.

The Mansfield Amendment (early 1970s and later omitted from defense appropriation acts) requiring the U.S. Defense Department to support only research with direct military relevance led to an emphasis on short range projects. While the pre-Mansfield projects of one major U.S. institution are still much referred to, their post-Mansfield projects have sunk without a trace. I don't suppose the Lighthill Report did much harm except to the British competitive position.

Government officials today tend to ignore science in planning the pursuit of competitive technological advantage. Both the Alvey and the Esprit projects exemplify this; DARPA has been somewhat more enlightened from time to time. It's hard to tell about ICOT, but they have been getting better recently. Some of the goals they set for themselves in 1980 to accomplish by 1992 require conceptual advances in AI that couldn't be scheduled with any amount of money. My 1983 paper "Some Expert Systems Need Common Sense" discussed this.

At present there is a limited but useful AI technology good enough for carefully selected applications, but many of the technological objectives people have set themselves even in the short range require further conceptual advances. I'll bet that the expert systems of 2010 will owe little to the applied projects of the 1980s and 1990s.

References:

Ginsberg, M. (ed.) (1987): *Readings in Nonmonotonic Reasoning*, Morgan-Kaufmann, Los Altos, CA, 481 p.

McCarthy, John (1983): “Some Expert Systems Need Common Sense”, in *Computer Culture: The Scientific, Intellectual and Social Impact of the Computer*, Heinz Pagels, ed. vol. 426, Annals of the New York Academy of Sciences.