wholly different circumstances): and yet it moves. From the Deep Thought chess program, to the MacnTax discs (do not laugh till you've tried to do US tax returns with and without them), to the new Carnegie Mellon speaker-independent speech recognizer, to machine translations getting steadily better, to expert systems and style correctors and the bits of vision and robotics on sale everywhere, we tend to forget how it moves, partly because so much does not come from approved AI laboratories. But that is another matter entirely!

Another disappointment in Lam's paper which makes it not really a worthy successor to Lighthill's, is that he does not know enough of what is going on in AI and so falls back on the tinkering of the gifted Whitehall amateur. Of course that is *ad hominem* and unfair, but he writes "The breakthrough would come if future translation systems can rely to a greater extent on the concept of probing a deep structure, rather than on a combination of morphological recognition, parsing—i.e. identification of syntax—and semantic methods—but this is far away". That sounds technical but in fact isn't: anyone who knows even a little of the field of natural language processing will know that the quotation is subtly muddled in more than one way (e.g. what could that "deep structure" possibly be in opposition to the all things he opposes it to?). One *could* give sense to the opposition, but nothing Lam writes makes one think that he has. This is not mere nitpicking, since some better level of understanding is vital to a serious discussion of this issue. Lam is also hopelessly Alvey/UK centred in what is now a genuinely world-wide field: he clearly does not know of the Carnegie Mellon speech work I mentioned above, and if he thinks there is no working machine vision, well! I certainly would not take any bets against the ATR speech telephony work at Kyoto based on what Lam says about it!

Some of Lam's discussion is valuable: it is fun to see some tally of where, at the edges of his (in my view wholly flawed) classification, Lighthill has been confirmed or otherwise by subsequent developments. He was right, for example, to be against general problem solvers. Those details alone make the paper worth writing and reading, but my fundamental complaints remain: the author's knowledge and analysis are superficial and that deprives the paper of real weight. Consider, in closing the following: "Diagnostic expert systems would be tiresome and expensive to program and to use—and indeed would lack justification—if they did not incorporate rules and hints provided by a human brain, which enable the system to economize on the volume of search and/or processing". The same muddles are all in there for such systems either work or they do not, they need no other justification. Of course any hints such a system has come from a human brain, where else, but that has nothing to do with whether they *model* a brain (which is what Lam intends but is unknowable in any case). Again, if he means that automation systems of this class need heuristics that express certain human shortcuts (as opposed to "complete" algorithmic solutions) he is just wrong: some do and some do not, and in the end only the market decides what is worth programming, making, selling or using.

Lessons from the Lighthill Flap by John McCarthy, Stanford University.

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 bureaucratic background of the author—or maybe they are 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 was not accepted then and has not become accepted since, because it almost entirely omitted the scientific basis of AI.

AI was not developed as a branch of biology, based on either neurophysiological or psychological observation, experiment and theory. It also is not primarily engineering, although an engineering offshoot has recently been developed. Instead it has been developed as a branch of applied mathematics and computer science. It has been used to study the problem of systems that solve problems and achieve goals in complex informatic situations, especially the *common sense informatic situation*. It has been employed in experiments and theories involving 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 does not. 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 argues 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 did not give up and have made considerable advances since then. I was disappointed when the BBC left that exchange out of the programme, 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 understanding 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 to be identified and solved.

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 US 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 US 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 is 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 could not be scheduled with any amount of money. My 1983 paper "Some Expert Systems Need Common Sense" provided a discussion of 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.* Los Altos, CA: Morgan-Kaufmann.

McCarthy, John, 1983. "Some expert systems need common sense" In: Heinz Pagels (ed.) *Computer Culture:*

The Scientific, Intellectual and Social Impact of the Computer. Vol. 426. New York: Annals of the New York Academy of Sciences.

Re Lam: Lighthill 17 years on by Karen Sparck Jones, Computer Laboratory, University of Cambridge.

There is an important point which does not get enough attention in papers like this on progress (or the lack of it) in AI. This is the size of the workforce. In Cambridge, for example, there are around seven times as many registered graduate students in biology as there are in computer science and, more significantly, as hardly more than a quarter of the computer scientists are in AI, there are about twenty-five times as many graduate students in biology as there are in AI. The absolute number of students in AI is not large either. This is of course not to claim that the pattern is the same in all universities and research training establishments. But there is no doubt whatever that for each body learning how to hack inferences there are very many more learning, under some biological or chemical label, to hack molecules. Moreover even if the comparison with biology as a really broadly-defined subject seems unfair (though computer science implicitly claims the same pervasiveness as molecular science), a comparison with physics, a reportedly static subject, shows there are nearly twice as many graduate students in physics as in computer science, and more than six times as many as in AI.

There is also the point that computer science is a postwar development, largely, along with AI, one of the last twenty-five years. Biological research as a recognizable base for modern biology has been going since the seventeenth century, and was thoroughly established, both intellectually and organizationally, in the nineteenth century. It takes a long time to build up a subject to supply the research cadre required to make and, even more importantly, to consolidate, intellectual advances. Physics papers based on work using facilities like CERN, for example, come with more than a hundred authors: teams and teamwork on this scale require the long-term subject preparation which physics, like biology but unlike computer science, has had. When compared with other areas of science AI, and even computer science, is still at a disadvantage, and it is therefore perhaps less depressing that it has made so little progress as surprising that it has made as much as it has.

A Rejoinder by Martin Lam.

I ought to feel diffident about responding further, since I am now in some danger of defending *my* views rather than those of Lighthill—which was not the original idea at all. However it is philosophically true, conveniently, that some propositions are untestable—for example, whether AI would have thrived without Lighthill; or whether if we only wait long enough we shall harvest rich fruits from AI. So I offer a commentary on chosen themes (which means that some points have to go unanswered—possibly, but not necessarily, because they are unanswerable.)

(1) Too much about the British scene; viewpoint typical of officials; not fair to have an outsider adjudicate. Yes, it is true that the British go in for outsiders. A currently popular philosopher-king is active in advising on takeovers and mergers, compensation for people jailed in error, the treatment of a former security officer and the freedom of the Press. So, for better or worse, the designation of Lighthill was in line with British practice—but well within the upper quartile of such appointments, since, given his background or backgrounds, he is surely very much an inside–outsider?