

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?

I admit that I had not succeeded in culling information about developments in the USA; a fault which the British AI sources must share. Yet, to be really nasty, I now try to make amends by calling in evidence an article in the Wall Street Journal for July 6–7th 1990 headed “Artificial Intelligence Industry hits Snag. Bright Outlook in US yields to Slow Growth, Big Cutbacks”. The staff reporter concerned records the dashing of the expectations of five years ago. Some market researchers had then predicted \$4 billion in annual AI sales by 1990; estimates are now for no more than hundreds of millions of dollars. “We’re in AI winter” says an AI newsletter—co-incidentally the phrase used by some protagonists in Britain to describe the post-Lighthill era. Another commentator quoted compares the AI boom to the tulip obsession of the 17th century. The article records the painful steps taken by US companies to put things right. Of course, I hasten to add, the AI community was not responsible for the wild estimates; however, one of the troubles, apparently, was that some “executives came from academic backgrounds and gave short shrift to marketing”.

My excuse for this bit of *méchanceté* is its theme, i.e. that the market has not, it seems, taken the view that AI has done fine in the USA—thus contradicting any argument that marvellous things are happening there, whereas we (UK) are still under a blight.

(2) More generally, there is the question how anyone can have the nerve to write about AI if he does not know this and has not taken account of that. OK, but is this not simply the converse of the thesis that only the insiders can be trusted to come to a view? This proposition, however, leads to a fundamental problem, if public money is to be committed. (Digression; this is without prejudice to the admirable thesis that more should come from industry, but it does not, here at least, in sufficient volume. Perhaps industry is guilty of short-termism; if so, this may reflect our way of using our cherished sovereignty. I understand that in a former period of rampant inflation the Head of a great technological firm required his managers to show an expected return of 27.5% before he would look at an R&D programme—in other words the “real” rate of return added to inflation. Is this not a rationale for being interested in applications, including those of AI?)

So—how is government money to be committed? Obviously where a “public good” has been identified, and, if industry is supposed to be interested in the shorter term then the State should be thinking about a longer period. (Like 500 years?) But how do you do this? Officials must come into it somewhere in the mediation of the answer; who else, if and when the “peers” are at a loss—surely not the very insiders who are arguing their case? Indeed, we must surely all know of instances (not, of course, in our own domain) where the interests in favour of a given programme have been close to fraudulent in their advocacy of their own interest. “Evaluation”—a current fetish—has even been conducted by some of the insiders who had partaken of the nosh. One reason for this has been the difficulty of finding anybody who had not either been one of the cooks or who did not aspire to be. This comment is not scurrilous, but points to a real live dilemma facing those trying to evaluate, say, European programmes. The presentation of projects, before, during and after, has become almost an art form, in which public relations techniques play a large part. It is Michie, who remains at odds with Lighthill, but who quotes, with approval, the conclusion of a meeting that “no-one gives objective scientific and technological advice”. Yet someone, somewhere, has to act on such advice.

In this light officials were probably glad to embrace the principle adumbrated by Lord Rothschild in the early 1970s that where research in government laboratories was not done for a specific real customer (like the Ministry of Defence before Gorbachov) there should be, for projects near to the market, a surrogate customer constituted by a Research Requirements Board. A generation of officials was nurtured on this doctrine.

(3) The definition of AI. This can be made to beg the question either way. If AI is a “task”—in that it seeks the means of carrying out functions otherwise executed by the human intellect—then do you include in AI self-focusing cameras or phase-locked loop tuning of FM circuits? If so, then at the limit all automation becomes AI. Surely there must be a recognizable anthropomorphic element? A trivial analogy that occurs to me—and for which I apologise—is that the dishwasher is a

very useful piece of equipment—but owes little to the way women, or even men, do the washing-up and is never claimed as part of AI. Conversely, I have seen an accounting package written in Prolog—a *tour de force*, but indistinguishable, for the user, from a similar package written in Cobol. Is this AI? I persist in my heresy (the prerogative of Galileo as invoked by Wilks) that one can try to identify as characteristic of AI, or in the contrary case note the absence of, the input of “intelligence” in some recognizable form that “models the brain”. (Wilks’ gloss.) This is the criterion I am trying to apply—not “whether practical applications yet exist, or which fields profit from them” (Michie).

What about AI as the union of applied maths and computer science? But this excludes, surely, some of the neuroscience? Then again, alas, poor Yorick shares my suspicions of cognition, while to Michie AI has an “intrinsic role as a vital bridge between the technology of computation and the biology of cognition”.

(4) Lighthill’s categories, and my re-using them, are to some degree arbitrary if not self-fulfilling. There is, I am bound to say, something in this—but only incidentally. For example he evidently regarded the work on language as coming under Category A, and I think it is logical to see the tree of development as mainly mathematical—Boole, De Morgan, Whitehead, Russell, Frege. I myself see something in the observation that the classification served not only the needs of taxonomy but mapped only too well on to the judgements. However, I do not see that Peter Jackson’s argument in itself necessarily invalidates logically a thesis that A and C are the worthiest categories; if, as he says, languages came from A with (let us say) a possible admixture of something else we might call C (though as I read Lighthill he meant by C, narrowly, research into the CNS) then the child does not have to be classified to B. Any general argument that an activity partaking of not A has to be seen as a bridge-building activity would incur the reproach of “fiddling” the classifications. Lighthill put DENDRAL in A presumably because that is how he saw it. More generally I see no reason why Lighthill, with so many strings to his bow, should find any particular domain—or a fortiori interdisciplinary areas—uncongenial and/or try to mould them to his own perception, as suggested by McCarthy.

(5) The weight, or lightness, of the resources in terms of student or research numbers attracted to AI. True, they are not overwhelming—but can one deny that molecular biology has attracted so many followers because it has indeed delivered a great many goods? (Incidentally, I defy any counter-argument that it is the Lighthill winter which has inhibited the luxuriation and the delivery of goods which might otherwise have ensued in AI.)

(6) Errors and omissions. OK, neural nets may be more useful than I say, (apparently even for sniffing) though I still find it hard to believe that an empirical device can do all the things claimed for it. An interview with a practitioner in the latest Expert Systems newsletter is far more sober in assessing the role of NN’s—as Croall calls them. He also prefers to see them as mathematically determinable rather than as working like a black box or an opaque bit of nerve tissue. In return a trivial jibe—what will come of the competition between the rival (or the same?) nets in the hands of the various institutions all trying to offer the safest but most forthcoming credit on the market-place? As for the importance of non-monotonic reasoning it is, of course, something we use a lot in our everyday life, but formalized or not it does not so far inform many AI programs—once again it is surely a means and not an end. I agree about the importance, as a tool, of RETE—even if it is now somewhat overtaken by TREAT as a matching algorithm. As for common sense, most practitioners, far from making it their aim to incorporate it, seem to despair of doing so except at the price of rigorous codification of rules determining each and every commonsensical phenomenon.

Machine learning—OK, so this is one more part of AI which is making progress. “Deep structures” to which I refer, is admittedly ambiguous; I meant the ones Chomsky was on about in the 1950s but I recognize that the phrase is also used to describe the (less) deep structures into which superficial aspects of language can be transformed in the process of machine translation.

Of course a lot is going on, but—to take back something of what I have conceded—if there are gaps in the appreciation non-AI chaps have of AI they will not necessarily be filled convincingly by a contents list of the topics that can be studied, however intelligently, by specialists, until these impinge on everyday life—or at least the marketplace, as invoked by Wilks. At the risk of sermonizing, I call attention to the danger that the AI fraternity may, as it were, tend to write for the aficionados and conoscenti—i.e. one another—and to discount doubts not only from the likes of Lam, who admits to being peripheral, but from the Lighthills of this world, whose credentials are surely (to use an ill-omened word) unassailable.

Looking back on my paper and the energetic but reasonably good-humoured comments of leading exponents of AI I do not regret my temerity in having a go at a (second-hand) assessment of AI and I am glad the Editor was willing to let the thing run. I recall being enormously impressed by the appearance on the scene of AI as a new approach to computing, and, since then, becoming somewhat discouraged; there must be others of whom this is true. I reckon that lurking behind the words and the argumentation there is a real imprecision about the content and objectives of AI, and that this is not just incidental but hinders progress. Perhaps I realise this more now than when I started writing. If officials have a faculty it ought to be that of dimly recognizing what is useful, even if the judgements are hard to externalize because they are “subliminally encoded”.