The Knowledge Engineering Review, No. 5:4, 1990, 265–276

Lighthill 17 years on

MARTIN LAM
Wembley, Middlesex, UK

Abstract

It is a common Establishment practice in the UK to ask an expert, recognized for his/her sound views, to give a view on controversial matters of public interest. There are recent examples in such varied fields as accounting standards, export credits, and microgravity. Administrations—for all the current fashion for "built-in evaluation"—have no time or inclination to take a look backwards at these autographs. In this paper the author, who was a civil servant, and has worked for ten years as a consultant, including two years in the European Commission, considers, with the help of hindsight, the Lighthill Report of 1972 on AI.

This review implies making judgements about the progress of AI which cannot help being contentious. First, there will be disagreements about judgements, or even assertions—the author cannot claim authority over the whole field, unless as a voyeur; there may be important developments he has failed to take account of. Second, there will be arguments that Lighthill was to an extent self-fulfilling—that is, that there would have been more advance if some aspects of AI had not been checked at his advice. Third, it will be said that it takes time to work up a new computational concept—or any other new thing.

However, whatever the reservations, the main conclusion put forward is that there was an element of over-optimism about AI in general which Lighthill had reason to castigate. In directing his comments to the utility of AI in specific fields Lighthill got some things wrong, but, on balance—which is not easy to strike—probably more things right than wrong.

Extrapolating to the present day the writer questions whether, in the UK after Alvey, and not least in Japan, resources are not being expended on AI objectives which are intellectually stimulating but remain too far from the application's sharp end; and therefore whether there is still not too much "jam tomorrow" about AI. Whether or not this questioning is judged to be wrongheaded, perhaps it will prove useful to have stirred discussion on the subject.

1 Introduction

The writer attempts here to assess, with liberal use of hindsight, the survey of Artificial Intelligence made in 1972 by Sir James Lighthill. The UK's Science Research Council (as it then was) sought an independent and personal view of AI by which to judge the increasing volume of applications for support received by the Council in this field. It could be instructive to see how far Lighthill's judgements can be seen to have been justified in the light of progress in AI since 1972. It is also relevant to the practice of forecasting to judge how far any individual, even a gifted scientist, can fulfill the function of official haruspex in a contentious environment. (There is an element of recursion here; the present writer is attempting a similar task, but with lesser credentials.) Last, but not least, it may be useful to conflate Sir James Lighthill's survey with the latest judgements about the Japanese Fifth Generation Programme and about the UK's Alvey Programme. The Lighthill text considered here, incidentally, is the version made available to the public, and does not therefore contain references to individual research projects or programmes current at the time.

A short digression may be helpful to explain how the effect of a strong opinion like that expressed by Lighthill can work through the system. The present writer was at one time an official member of a Research Requirements Board and, coincidentally, was called for jury service during

the same period. The workings of the jury gave an insight into the group psychology of the Requirements Board. The latter—like the jury—was composed of mature persons endowed with goodwill and honest intentions. But when presented with a proposition for discussion they did not spontaneously come to a view. On the contrary, there was a collective reticence which at first was disconcerting. The jury behaved similarly, and were only too grateful to any juror who would come out with a strong view which would help the others to reach a conclusion. (For example, "That policeman was a liar".) And so with the Requirements Board; an opinion offered by an individual member was likely to have its contrast enhanced and then form the basis of a collective judgement. This was not a conscious process, but rather the only possible way of reaching a consensus when the members of the jury, or the committee, were genuinely at a loss.
This analogy may help to explain how it was that the effect of the Lighthill Report, particularly for the great majority who had not read it, was far more potent than its literal interpretation would now seem to have justified. The reputation of the Report as generally unfavourable to AI made possible a balancing of collective perceptions in the scientific establishment about AI in general towards a flop rather than a flip when, without the Report, there would have been a continuing state of hunting and somewhat random decision. The Report cast something of a shadow over AI as a unitary research area during the 1970s; only in the early 1980s was there a come-back, mainly as a result of the Japanese Fifth Generation programme.
In his Survey Lighthill distinguished between work on Advanced Automation (his Category A) and Computer-based Studies of the Central Nervous System (Category C), each of which he expected to develop along its own path, using methods appropriate to the specific field—drawing on engineering and mathematics in A and on psychology and biological sciences in C. He heavily discounted Category B (for Bridges) and predicted that far from developing so as to unify the whole area of study (as some expected at the time) Category B would wither away. This reflects his view that progress would be made by concentrating on specialized fields with, in each, a modest "universe of discourse", rather than more general studies. We look at the discrete disciplines or sub-disciplines reviewed by Sir James Lighthill and compare his readings with those we can make after seventeen years. In doing so we can still follow his taxonomy of categories A, B and C, which was essential to the analysis, though perhaps by now it tends to constrain the order and the range of topics discussed. (The present writer, attracted by the technique, has added a few more alliterative headings—for which he alone is responsible.)
It would be customary to offer, "up-front", a definition of Artificial Intelligence, to forstall those who are tempted to say that "it all depends what you mean by . . .". Perversely, the writer prefers to leave definition to nearer the end; different perceptions, over time, of what AI is about may turn out to be the key to differing views about AI, both as to its utility at a given moment and to its prospects as an enabling technology or set of techniques for the future.

2 Lighthill in "micro"

2.1 "A" for advanced automation

Among the first applications commented on by Lighthill within his Category A ("Advanced Automation") is pattern recognition in the context of *printed and typewritten characters*. He identifies the industrially important aim of machine recognition of print or typescript. He comments that "in this area good progress had been made . . . in spite of a quite surprising degree of difficulty". This is pattern recognition of a specialized kind, and pattern recognition is something AI has tried hard to master, since the eye and the brain do it very well. There is also an element of machine learning, in that the system has to be coached in the recognition of letters it tends to get wrong.

Lighthill was right that recognition was being accomplished despite the difficulties. Indeed by now there has been reasonable success, charted by the sale of marketable products, in systems which scan a page of print or typescript, and by recognizing the patterns in a matrix translate the bitmap into ASCII code capable of subsequent storage and word-processing. These systems are not perfect; as with machine translation systems the text needs some editing. This can be time-consuming, and so can the use by the operator of the provision to "teach" the machine to recognize better some of the characters it gets wrong. Given clear object material the whole process is now good enough to enable whole text searches to be made on material originally scanned on to a CD-ROM. However, the point of record is that Lighthill rightly identified the recognition of printed or typewritten characters as an area where good returns to investment in R&D could be expected. (Though, to be fair, it should be added that though the recognition of patterns remains a classical AI objective the algorithms used here in the scanning program need not owe much or anything to AI as it is generally perceived. Here, as in other cases, computing emulates what human intelligence does, but works in quite a different way to do the same job.)

As regards the recognition of hand-written characters Lighthill comments that it "appears completely out of reach". This has been mainly true, with signature verification as a shining exception; the recognition of postal codes may become another. The Japanese have systems to recognize clearly handwritten kanji and kana.

Lighthill is almost as hard on *speech recognition*. This had been "successful only within the confines of a very limited vocabulary, and large expenditure on schemes to produce machine recognition of ordinary speech has been wholly wasted". Learning techniques might be effective for "an exceedingly small vocabulary . . . like the decimal digits". Surely this judgement is still valid in that no system for the recognition of continuous speech has yet hit the market? Very reputable companies have demonstrated to visitors prototype systems which accept measured dictation and turn it into text. They are said to have a vocabulary of 2000 words. Systems based on the recognition of spoken numbers are on offer commercially. The speech processing part of the Alvey Programme contained six projects, including the development of algorithms and the recognition of connected speech—including Chinese for word processing. No commercial exploitation seems to have yet been announced, though several developments were to fructify by 1989.

Whether all this is success or failure for AI is a matter of calibration, as with Dr Johnson's classic reference to a biped dog. One can be amazed that recognition, based, as Lighthill predicted, on semantic as well as parsing networks, and significantly, on the signal processing of phonemes and the Fourier analysis of their frequency distribution—not directly attributable to AI—can be so effective. Artificial intelligence makes a contribution by the use of rule-based systems in the recognition of the phonemes and in parsing. Yet one can still have a strong feeling that operational use is far away—even though both the Japanese NTT and British Telecom predict, apparently seriously, a future on-line oral translation facility whereby two businessmen using different languages could talk to one another on the telephone—piling the Pelion of on-line translation upon the Ossa of continuous speech recognition.

About *machine translation* of text Lighthill thunders that it is in this area that "the most notorious disappointments have appeared". He infers that, as with the use of natural language in communicating with a computer, translation can succeed "only where a very detailed knowledge of the universe of discourse is stored within the machine, . . ." This augured "badly for the future availability of machine-translation programmes versatile enough to be commercially viable." In the event several systems are commercially available; one of them, Systran, is being developed continuously in the European Commission for pairs of Community languages. Supported by growing context dictionaries and owing little to AI as generally recognized, it does work, in the sense that the meaning of official documents is conveyed reasonably enough, and in a style which is not always worse than the original. However, as with all translation systems, including those sold widely in Japan, the text has to be edited to bring it up to idiomatic standard and to put right occasional misunderstandings. (As an unfair example, "satellite channel" was once translated for the writer as "Manche de satellite".)

As before, one of two views can be taken—either that it is a marvel that Systran works as well as it does, or, alternatively, that many Systranslations, based on ever larger linguistic databases, though superficially accurate, lack just enough of idiom and style to be off-putting and, incidentally, to make editing that much more difficult because the texts are, in a sense, accurate without being idiomatic. The practitioners can retort, not without reason, that some apparent shortcomings reflect the turgidity of much Brussels-speak.

Japanese systems translating into European languages are likewise a *tour de force*, and the need for editing at the 10% margin does not impair their acceptance by buyers, who apparently still make a net gain of time and resources. Alvey had a go at a package for technical translation into and from Japanese. As with many automated systems machine translation may be more useful as a help and support to the human agent than as his/her replacement—if only the professional translators would see it that way and learn how to write and think on the screen. A second generation system—Eurotra—is in the EEC pipeline as a long-term research project. It will work not by language pairs but by "translating" each language into a representation internal to the machine. A transfer process will then put this representation into the target language. The AI in Eurotra is mainly concerned with parsing and with the rules which regulate the transfer process. The justification for support—if research funding is considered, realistically enough, on a bookmaking basis—lies in the long odds. That is, the prospects for success may be modest, but the stake is justified by the huge economies to be reaped if machine translation is in time operationally successful.

Machine translation has in fact come out better than Lighthill expected, though so far largely with the help of conventional computing. Systran copes—at the expense of its context dictionaries—with the generalized administrative "mandarin" in which the bulk of Commission papers are written. Specialized material will require the building of more context distionaries, but Eurotra will be aimed from the outset at several different types of text.

About *intelligent data storage and retrieval systems* Lighthill observes that "quite good performance has been achieved where the problems were confined to data of very precisely defined and analyzed types. By contrast, generalized information-retrieval systems have been somewhat disappointing . . .". This judgement remains valid today, in spite of Intellect, Atrans and the Nattie project, in the sense that the AI techniques of intelligent accessing of databases and/or the use of natural language work best in narrow domains in which the machine can reduce a question to its essential elements. The combinatorial explosion so far makes it uneconomic in time and processing to traverse the network iteratively until all ways of putting a question can be reduced to a standard form. Alternatively, if the user agrees to phrase his questions in a given form, and/or cares to store his information in a way which is conducive to retrieval by a generalized formula, he can put his queries in near-English. Otherwise there has been little or no progress in any generalized paradigm for the ordering of information in such a way as to provide a technique of access more like the ways used by the human memory. The real progress in retrieval, referred to earlier, has come from the brute force whole text search based on key words, made possible by faster processing of the conventional matching of strings.

More recently the hypertext approach to the ordering of information could be argued to come close to the free association which characterizes the human mind— but this has few demonstrable AI connections, unless indirectly in the use of lists. Topic Identification is a promising hybrid approach. If pattern matching via genetic—and therefore learning—algorithms is made to work, by mediating a competition for survival between the results of successive searches, then this will be an AI contribution to retrieval. Object-oriented languages have been adapted to a mapping process which can make storage and retrieval from relational databases more flexible in a package available commercially; however, AI is only one field to profit by, and thus to be given credit for, the introduction of object-oriented languages. The USA is reported to be considering the use of OO databases for genome information—while the Japanese are using for this the database which is part of the Fifth Generation Project.

2.2 "C" for central nervous system

Lighthill's next main category of applications is labelled C—Computer-based CNS (Central Nervous System) Studies. He comments that "in spite of a respectable volume of achievement

. . . most workers who entered the field around ten years ago confess that they then felt a degree of 'naive' optimism . . . . It is, once more, the most generalized types of studies whose end-products have proved most disappointing". Lighthill adds that psychology has gained more than computing from the work in this field—in other words that experimental psychologists who had a flair for, and have taken the trouble to learn about the working of the computer have found the building of computer models helpful to their understanding of possible configurations of neural paths, the propagation of signals and their processing in the brain.

Lighthill identifies in particular the study of neural networks as a case in point. He could not guess, however, that in the late 1980s the converse effect would be observed, and that a branch of Artificial Intelligence would be devoted to configuring a number of processors in so-called neural nets with the objects of matching patterns by imitating the way signals are believed to be processed in the nervous system. So the tables have been turned; the psychologists have little hope of getting nearer to the simulation of neural nets through modelling, because of the multiplicity of connections between synapses. It is the computer scientists who are turning neural networks to advantage by emulating them (as best they can) on a very simplified basis. The salient characteristic of these man-made networks is that the processing and the communications functions are combined; the connections are dynamically given a weighting corresponding to the importance they earn as the computation proceeds. The system "learns"; for example, a historical (past) series can be fed as signals into the inputs; the machine is programmed to change the weightings used in the matrix multiplication of these signals, and/or the subsequent computations based on them, so as to give the historically valid output. With luck the system can then be used predictively. In the same way it can recognize frequently-occurring patterns.

The results achieved experimentally in Europe, USA and Japan have led to the offer for sale of software which emulates neural networks for trial or use in specialized applications. Intel is selling a chip for such constructions. Lighthill did not predict these developments, but he will not be surprised first, that the proposed applications are relatively narrow, and, second, that no experimental machine has coped with anything like the number of neural connections which operate even in the most lowly vermiform invertebrate nervous system.

Using a similar judgement that in liaisons between computing and other disciplines the latter were likely to gain more than the former, Lighthill noted the relevance to computing science of what he calls psycholinguistics, but suggests that in the event psychology—rather than computing—would profit from the synergy. The psycholinguists would learn from the new algorithms, based on semantic as well as grammatical sorting, how the brain understands language. In practice the exchange has turned out to be two-way, though, as pointed out above in the context of machine translation, AI has not yet mediated a real breakthrough. A number of stages are required, few of which are AI-related. The breakthrough would come if future translation systems can rely to a greater extent on the concept of probing into 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.

Educational psychology receives a significant comment from Lighthill. He sees, rightly, that Computer Aided Instruction of the first generation was in need of greater flexibility and a facility for targeting, and that Artificial Intelligence could in principle be harnessed to this end. He saw that branching programs were by no means the end of the road. However, Lighthill then goes on to comment on the possible generation of new teaching material (rather than that which is simply pre-stored) and judges these "applied research dreams" to be "singularly untimely and unpromising". It is clearly still true that teaching programs have to be based on material already stored and that any idea of programs which would produce new teaching material using natural language remains unfulfilled. It was at this imputed objective that Lighthill's strictures were aimed, though, truth to tell, it is not easy to find practitioners who admit to having believed in it—though there have been some experimental attempts at extracting material not stored in finished form as such but synthesized by mapping the subject matter on to information patterns stored in the machine. What has come about in practice in this field is rather some modest support from AI, in the form of progress in the intelligent modelling of the learner—i.e. in ways of identifying the hangups and propensities of the learner by inference from his/her responses, in monitoring the learner's progress and in devising a course of learning to suit the individual.

2.3 "B" for bridge building

It is a feature of Lighthill's survey that his classification of his subject matter mapped elegantly and conveniently on to his broad conclusions; having identified some high spots in Categories A and C, he finds a valley of under-achievement in his Category B—defined as Bridge activity or as Building Robots. On the grounds that it is this category which, as it were, claimed to bind the whole discipline together Lighthill questions whether "the whole concept of AI as an integrated field of research is a valid one". This is a crucial question to which we return later.

As one of the reference points by which success or failure must be judged Lighthill quotes, not unfairly, comments by unnamed scientists (they were "able and respected"—which made it worse) that AI represented "another step in the general process of evolution". We discuss below the extent to which AI, or some aspects of it, have been subject to some degree of what we should now call "hype". Suffice to say here that Lighthill was obviously put off by the "wild predictions"—though still falling short of Rossum (Capek)—being made about the future of generalized robot development, which he saw as a large component of his Category B.

In his specific criticisms aimed at the allegedly *intelligent robot* he discounts in particular the work on vision. Surely machine vision in the wide sense is still out of reach? Engineering companies have a form of it as an option, for example in identifying where a screw hole is in processes involving automated picking, inserting and assembly. But, in spite of Alvey's efforts in image processing, is there current operational use of vision with perspective, possibly in three dimensions and possibly with the help of radar or sonar techniques?

The invaluable services rendered by the machines which spray paint or weld as instructed by the human operator, and which, as obedient robots, rely as much on the emulation of muscles and joints as on intelligence of an imitative kind, are consistent with the Lighthill thesis that narrow advances are more likely than wide breakthroughs. Similarly the effective use of numerically controlled machine tools and more responsive control systems can be seen as practical steps towards robotization. Whether these systems involve AI may be a matter of philosophy, but in any case Lighthill would classify them to his Category A—"automating the solution of problems . . . in practical fields of application". These both work and are useful; this he evidently could not say for the abstract robotics research he had seen. (However, he might be surprised by the French project for "Robots automoteurs multiservices" designed, as the name implies, to develop a series of robots adaptable to a variety of tasks.)

3 Lighthill in "macro"

3.1 "F" for future

Following his judgement by categories Lighthill goes on to offer a view of the "next twenty-five years", modestly foregoing any speculations going beyond the end of the century. His macro forecast is that there will continue to be a "bimodal" distribution of achievement—i.e. within Categories A and C—with the valley between them—B—even flatter than it was in 1972. Applications will be the strongest runner, based on increasingly firm links between "each research area and its field of application". As examples he cites techniques for advanced automation devised by researchers who have practical knowledge of the domain and are also in contact with control engineers. This process would enable the development of "advanced computational theory", including new programming languages. Certainly this is the path by which CIM and Flexible Manufacture have indeed advanced, as noted below.

Lighthill refers, but only in passing, to two programming languages without naming them; LISP had been going for more than ten years when he wrote, but PROLOG had only just been launched. Relying on first order predicate logic PROLOG, developed so as to enable it to work concurrently, was to be the basis of the inference machine in the Japanese Fifth Generation Project. However, the undoubted ingenuity and elegance of PROLOG as a means of declaring knowledge in the form of propositions, which can then be logically resolved to yield an answer, can still fall down at the Lighthill hurdle of computational overload when a large number of axioms or statements has to be traversed. This difficulty can be reduced by grouping the rules (as with XCON below).

Lighthill quotes in the same context the automation of factory assembly processes—what is now called Flexible Manufacturing or Computer Integrated Manufacture. He considers that progress here will be connected with the evolution of CAD. (Only now is the use of expert systems in manufacture beginning.) Lighthill recurs again to his theme that specialization and fragmentation may afford a better path to success than research of a generalized nature, quoting theorem solving and cartography as successful cases in point. CIM and Flexible Manufacture would be successful instances of Lighthill's Category A.

Lighthill took very seriously the linkage between computing and work on the central vervous system; he sees that computing has something to learn from illumination of how the CNS works as well as the other way round. However he makes a characteristic judgement (in that he sees it as his duty to indicate boundaries where he perceives them) that "simulation of the extraordinary self-organizing capabilities evolved by the human CNS may actually be unattainable through ingenious software developments". In the same vein "With the resulting growth in understanding of the human CNS, respect for it may, perhaps, grow to an extent that will reduce the ebullience characteristic of past predictions of AI possibilities". Looking at machine vision on the one hand and our simulated neural nets on the other, can we disagree?

3.2 "G" general; "H" for hype; "J" for Japanese

After this review of the detailed comments made by Lighthill it is time to consider some of his general conclusions—both explicit and inferred. First, there is the castigation of "hype" (our modern term, not his; if favourite technologies are doomed to end up following a sigmoid path they often start off hyperbolically). "Most workers . . . confess to a pronounced feeling of disappointment in what has been achieved . . . . In no part of the field have the discoveries made so far produced the major impact that was then promised." "Claims and predictions . . . went even farther than the expectations of the majority of workers in the field . . . the lamentable failure of . . . inflated predictions."

Surely Lighthill was right in diagnosing a degree of euphoria on the part of some workers in this field? History has, to some extent, repeated itself. The Japanese Fifth Generation Project, launched in the early 1980s, was extremely well prepared, and has been pursued with zeal and a significant intellectual investment. Yet the last Conference (November/December 1988) confirmed that—at the most optimistic—it will take far longer to achieve most of the goals, and/or that—at the bleakest—some of the objectives will prove unattainable. (In fairness, it may be that the choice of logic programming as the core of the Project was a mistake—particularly with an eye to future parallel programming; yet this choice itself reflected a loyalty to A perception of AI.)

Some media treatment cannot, of course, be blamed on the AI fraternity. In any case the progress of technology is fascinating; can we be surprised that those who write and televise about it are enthused? Or that those in quest of funds for research emphasize the striking results at the end of the road rather than the tentative nature of their initial explorations?

Second, Lighthill's Leit-motiv that more AI mileage is to be got out of specialized work in a suitable field than from generalized attempts to devise all- or multi-purpose systems, is supported by work in a sub-discipline of AI of which Lighthill had no cognisance under this name—Expert Systems (though MYCIN was already extant and DENDRAL is mentioned with approval by Lighthill). Certainly expert systems are becoming a commercial success, in useful though often modest ways. The large ones that have become operational, like X-CON and its companion programs, which, after nearly ten years development time, help DEC to configure its computer systems, have been focused on particular fields or problems—just like the human experts.

This is not to denigrate the utility of the large and widely applicable commercial development systems which, if users can learn how to use them, are a real support to the writing of expert systems, nor that of the smaller and cheaper tools which do the same for less ambitious users. It is consistent with Lighthill's belief in progress based on a narrow universe of discourse that success has come from synergy of machine and expert, and not by the machine's superseding the man. The way expert systems have developed is, by definition, in line with Lighthill's preference for a specialist rather than a broad-brush approach—even if his judgement was given without an appreciation of the path to be taken by expert systems.

3.3 "E" for explosion

Ingenious knowledge-based systems have to some extent overtaken the *a priori* criticism made by Lighthill that combinatorial explosion would frustrate many attempts in computing to do things which the brain can do well, on the grounds that the volume of computation would be too large. His view would accord with the mathematical demonstration that some problems are computationally unsolvable, while others, though in principle solvable by the archetypal Turing machine, demand a solution time which increases exponentially with the number of inputs. The combinatorial explosion—resulting from the need to convolute the many combinations of real life possibilities—can frustrate the application of purely mathematical methods. Sometimes ingenious programming techniques can lead to striking improvements in speed. However, writing in 1972, though he recognizes the use of a heuristic approach, Lighthill seems to underestimate the potency of human input in the form of specialist knowledge and of logical short cuts in programming so as to achieve economies of search and processing—i.e. what is usually called heuristics (though this term would be more properly used to mean the process of discovery by the machine).

By a gradual process of redefinition, such engines are now counted as an important part of AI. Intelligent Knowledge-based Systems (IKBS) are, by definition, those which rely not solely on machine reasoning but, indispensably, on an input of humanly organized knowledge on the one hand and human reasoning on the other. Indeed IKBS maps well on to Lighthill's Category A. There are still constraints derived from the combinatorial explosion, but the boundaries have been pushed back by the AI practitioners themselves, who have devised new reasoning tools to yield short cuts. Moreover, by now there have been striking advances in the accessibility of high speed computation as its incremental cost has gone down. It is surprising how much computation is now acceptable, as in the case of the XCON and XSEL systems, approaching 20–30,000 rules and configuring more than 30,000 hardware and software components.

Thus AI as it is now presented offers, above all, the chance of doing certain things more economically by the use of short cuts which rely on basic human mental processes. 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 economise on the volume of search and/or of processing. Indeed one attempted definition of AI is "a technique for solving exponential problems in polynomial time". Such techniques as truth maintenance, whereby the system can test successive hypotheses for falsity, are derived from AI.

Lighthill sees clearly enough the scope for an input of human knowledge and skill, and makes the debating point that this impairs the pure artificial intelligence concept which was favoured at the time in some circles. "The widespread view that only heuristic methods are effective is a serious setback to the more extreme versions of the AI philosophy." By now most practitioners of AI would surely accept that the heuristic approach, supported by other inputs based on human understanding, is the path by which AI offers a significant complement to conventional computing. Few would now espouse the extreme view that AI is destined to surpass human capability, or even supersede human use of conventional calculation.

At this point the writer has perhaps to be careful of not conceding too much to AI. AI can provide techniques which reduce the amount or the area of search, but this does not imply that all reductions in search are derived from AI. Ingenious search algorithms have been devised, and are incorporated, for example in chess-playing programs, which enable them to engage depths of play which would have otherwise seemed impossible because of combinatorial explosion. This view may be contentious, but the mini-max technique, though derived from Shannon and Turing, is surely not part of AI and neither, surely, is Alpha-beta pruning, even though it is a way of narrowing the search and thus of avoiding the brute force exploration of possibilities. (Though, of course, in the chess context the identification of the variables which help to distinguish winning or losing configurations and the choice of the weightings to be given to them owes everything to an original human judgement.)

The specialized processors now used for chess equally represent a *tour de force*. More generally one can see with hindsight, what Lighthill could not know—that at a certain stage machines would be offered for symbolic computation in the sense that they would be tailored to the processing of lists stored with the help of pointers and descriptor tags and using microprogramming for the compilation of LISP functions. This is another reason—outside the scope of this note—why AI has done rather better than Lighthill predicted in coping with a large volume of computation.

3.4 "P" for plaudits

Although Lighthill did not know "Expert Systems" by this name he goes out of his way to praise DENDRAL—the famous, if not legendary, program enabling the user to infer chemical structure from, *inter alia,* mass-spectroscopic data. He regards it as a success, but one not necessarily attributable to AI, on the grounds that it is based not on a set of stored axioms or on propositions from which logical inferences may be drawn, but on the combining of probabilities supplied by human estimation. This must seem to us a somewhat astringent view; Lighthill is seen once more to be drawing a distinction between the "pure" AI then promoted by some practitioners and the "assisted" form in which it functions with, and not in place of, the human operator. Even in present-day terms DENDRAL is a jewel in the AI crown, having established techniques which are now in general use. The real surprise is that there have not been more DENDRALs in all these years, and that it is not clear how much even DENDRAL is being used in earnest (though DENDRAL, or son of DENDRAL, appears to be behind a current super-computer-based molecular design package).

In a postscript Lighthill also generously records how impressed he was with the PhD thesis by Terry Winograd "Understanding Natural Language", published in UK in 1972. He evidently admired, apart from the style, the success of the program, developed in the thesis, which enabled a conceptual world of blocks to be given instructions and conversed with by the use of a near-natural language—in a necessarily limited universe of discourse. Lighthill concedes that this came near to the building of a conceptual robot—and thus to a success in Category B. However, Lighthill credits the good marks earned by Winograd to Category C, as an exercise in psycholinguistics, rather than to Category B. In retrospect this was surely a correct judgement; there is no evidence that robotics was able to use Winograd, nor that the *tour de force* implicit in the technique could be applied to a wide world.

3.5 "R" for the reality of artificial intelligence

Was Lighthill right in questioning the existence of AI as a unitary discipline? The key to this theme of the Report is the gradual but noticeable change that has taken place in the understanding of what Artificial Intelligence is. The common-sense definition or denotation of AI remains in substance the same—the study of those aspects of computing which imitate or emulate the way the human brain is believed to work; this definition would include the use of computing as a tool for research into the workings of the central nervous system. Yet there has been a shift in the connotation of the term AI. People now regard themselves as working on AI as much by reference to the techniques they apply as to the subject matter of their endeavour, which they may reasonably regard as getting the computer to do things which in a human being would require intelligence and not merely infinite patience in sifting and calculating. It will have been seen as we traversed the Report that this shift of meaning leads, at best, to some misunderstandings, and, in several cases, to the need to revise some of the judgements made in 1972.

So in this light how should we now define AI? One approach is to consider as within the domain of AI all computing which deals with symbols rather than numbers or letters, and/or which relies on rules rather than on algorithms and/or which deals with probabilities rather than yesses and noes, or with a band of acceptable answers rather than uniquely valid solutions. A US guru has defined AI as a "collection of programs, ideas and tricks for processing symbolic information". One attribute typical of an AI approach would be the inclusion in a program of elements of heuristic (i.e. assisted) search based on human knowledge of the field. Where a problem can be solved effectively and economically by conventional computing there will be little or no point in using AI methods. Lighthill saw this clearly in his references to the effectiveness of Automatic Landing Systems and to the counting of features in microscope slides without the help of AI techniques. The latter more legitimately come into consideration when a problem is unmanageable if attacked by conventional methods. (Indeed, it is not a bad rule of thumb, used intuitively by many people, that you first see whether your problem can be economically solved by conventional computing, and only if it cannot should you consider AI.) If, however, the approach is rather that AI includes the use of any techniques which help the computer to solve problems which would require intelligence in a human being then, the writer concedes, you can incorporate game playing, search algorithms and almost any piece of ingenious programming. But, if so, how does one then distinguish AI?

3.6 "R" for wrapping it up

After weighing the judgements offered by Lighthill and a review of how the nature of the target has changed, as well as moved, since he drew his bow, we can briefly consider how Artificial Intelligence looks now. Definition or no definition the specialist libraries and bookshops are full of volumes about AI. Robotics are there in force, dealing both with what can be done right now and with the study of what might be possible in future—including much musing on machine vision, "planning" (handling by use of the computer the causal and temporal relation between events) and machine learning. The key to an effective solution to many problems in computing lies in finding the right "representation" of the problem. Many researchers are willing to help in this stimulating quest. AI has certainly burgeoned since Lighthill wrote; the buds and offshoots deal with aspects of AI that might well have seemed peripheral or of doubtful relevance 17 years ago. Expert systems is now a thick shoot—perhaps the only one bearing any immediately harvestable fruit. A lot of work—perhaps a disproportionate amount—is being done and studies published in support of expert systems. Some deal with the various ways of transferring the knowledge of the expert to the machine. There are manuals on statistically-based inference—including the aptly named fuzzy sets and variants of Bayes theorem combining the conditional probabilities of several events.

Importantly, much ingenuity—and public investment—is being devoted to the development of parallel processing, in the expectation that the big discontinuity of the future will come when sequential computing will be complemented by operating systems and programming languages allowing the machine to perform simultaneously, as the mind would, such operations as matching, which at the moment have to wait in a queue. Parallelism—up to a point—has been available for some time in the hardware sense of processors that can do the same operation simultaneously on different data, or different operations on identical data. The problem with more complicated programs has been to arrange the necessary communication between processors; operating systems for parallel working in support of languages providing concurrency are still primitive and, as the Japanese have found, both face a long development time. There is a gap, perhaps temporary, based not on credibility but rather on the difficulty of getting to grips with replacement programming concepts—as with Occam, which supports that insufficiently considered device the transputer. Concurrent searching should reduce by several dimensions the time required, say, for parsing of natural language. But here we go again—so much of the jam is for tomorrow—or will it rather be our future bread and butter?

Some of the academic work on the borders of AI relates to "cognition". This word was not in Lighthill, but is in common use now to denote the study of how the mind perceives the external world, is aware of it and learns about it. There is now a recognized place for cognition in the design of the interface between the person and the machine. However practitioners may take the present author to task for his reference to the "borders" of AI; some cognititians or cogniticians are believed to claim that AI is in truth part of their discipline—if, indeed, cognition covers everything connected with the workings of the senses and with understanding the signals they process.

3.7 Is it C for confutation?

All these areas of study are estimable, ingenious and intellectually demanding. They also have an arguable family coherence. Is their existence, and indeed their proliferation, a confutation of Lighthill? Insofar as what is coming out now is still groundwork—a consolidation of the foundations of AI, rather than a construction of applications—Lighthill is, within his own terms of reference, vindicated. He was searching essentially for the creneaux, the commanding heights, of the discipline, and those likely to lead to real life applications. Surely these cannot yet be confidently determined? Where he seems, in retrospect, to be less right (if this quantification is allowable) was in his unwillingness to recognize the identification of the whole field as possessing common attributes and therefore worthy of an integral research effort. As we have seen redefinition has largely overtaken this argument. OR HAS IT?

Clearly it is with possible applications in view that the US and Japanese multinationals are selectively active in AI research, though there must surely be more than a suspicion that if they were wholehearted about it there would be less need for public or collective industry money to be used. In Japan ICOT continues steadily; in the US DARPA and the MCC both have an AI component, and so does ESPRIT in Europe. Under Alvey in the UK a proportion of official funds was channelled into AI-related research, always with an application perceived, albeit sometimes distantly, at the end of the road. The Bide Report would have aimed even more at applications.

In the assessment of the Intelligent Knowledge-Based Systems part of Alvey the officially appointed reviewer—a kind of latter-day Lighthill—counts it as a gain that "the academic AI R&D community has been expanded by an estimated 150 or 200 professional staff" and also notes the "expansion of the industrial AI R&D community". This quote would be unfair if it were meant to imply that that is all the reviewer has found. However, it is not unfair to point out that most of the Review assesses the measures taken to revive AI; thus the results to date still amount to an accumulation of means rather than an achievement of end products. The reviewer rightly points out, of course, that Alvey was a scientific rather than industrial activity—the international trade rules prescribe that this should be so. Even so, is it not still fair to ask where the fruits are by which AI would like to be known?

At this point we can return to the questions of the opening paragraph. It can be seen that Lighthill was wrong in some of his judgements about why, and where, AI was not likely to meet expectations—but not so wrong in toto. Given the world-wide scale of work on AI we can rule out the possibility that the Survey was self-fulfilling, i.e. that buds have not blossomed because they were nipped by Lighthill—or, more properly, by the scientific establishment, which could have disagreed with him. As to the utility of the kind of survey done by Lighthill there are turning points, at which crucial judgements have to be made by the State apparatus. This happens both in the mixed economy generally and on the technological scene in particular, where governments, in pursuit of various public goods, spend more than companies. Since these controversies necessarily involve a push and pull between opposing interests it is reasonable that disinterested and therefore objective reviews of a problem field should find a place. Appointing an individual to arbitrate at least obviates internal controversy—he has to handle any incipient schizophrenia autonomously. What is rare is an audit of such reviews, forecasts and clairvoyancy in general.

As to how far AI can be known by its fruits, even if Lighthill went rather far in 1972 in reproaching AI (by implication) for its failure to come up with deliverables, is there not more than a suspicion that the pendulum may now have swung the other way? The intellectual and other resources deployed in AI research are surely valuable in terms of opportunity cost—particularly when surveys suggest that industry generally sees different needs from those favoured in universities. The review of the field in this article may be interpreted by some as a demonstration of how much promise AI has accumulated. Even so, is it quite fair that AI, now stronger still and wider, should continue to keep the user waiting? Or is it that everything new takes a long time to find application—as with lasers, holography, superconductivity? It is arguable that it was too early in 1972 for Lighthill, with all his insight, to identify successful applications, and that he was wrong to believe in only what he could touch. But "*quo usque tandem*" (as the man said)—for how long do we still have to suspend disbelief?