Part III Comments on Parts I and II

1 Comment by Dr. R. M. Needham

Computer Laboratory, University of Cambridge.

Like all classifications, Lighthill's division of AI into three main parts is contentious in detail as doubtless was Caesar's similar dissection of Gaul. It would not be useful to discuss whether particular individual activities are best placed in A, B or C – at any rate if one accepts as I do the spirit of the classification. Since I basically agree with Lighthill's conclusions there is perhaps less to say than in Sutherland's commination.

The aim of the category A work is technological. Any method which achieves the desired result will do, provided it is not too expensive. This is by no means to say that it need not be founded on detailed knowledge of the subject matter, nor that it should eschew devices such as the use of heuristic methods which are perhaps associated with AI rather than with automation. On the contrary, the use of reasonably reliable heuristics is very suitable to the "no holds barred" approach. Heuristics, in general, are devices to avoid excessive searching by acting on guesses as to where to look – guesses which are not provably correct but which usually lead to something sensible. For example "To find a letter from the SRC in the Departmental office, look in the file marked 'SRC'", or "To proceed from Cambridge to Edinburgh, go first to London". Neither of these is always reliable, but both are based on a sound knowledge of the relevant facts.

It could perhaps be said that in any knowledge-based system some of the access rules will be formalisable and embodied in regular programs, and some will not and will thus have to be treated as heuristics. At this point, a question arises which Lighthill does not treat very much. Are there any general principles – that is, principles which apply to numerous applications – which guide or might guide the application of heuristics? Workers whom Lighthill might describe (? stigmatise) as being in category B say that such principles are, inter alia, what they are looking for. Some of the justification for this kind of work would seem stronger to those who believe that such general principles are there to be found. Lighthill suspects that they are not, so it should be pointed out that the Hart-Nilsson-Raphael theorem is one such that has been found. I do not personally think there is much to dig for here, but one should not deny that there is anything at all.

Lighthill's category C is quite outside my own technical knowledge. It is self-evident that enquiries, computer-based or otherwise, into how people work constitute an important field of scientific endeavour.

The present question is about the intelligent behaviour of people, or the way people function when behaving intelligently, and whether any work can be important to this which does not explicitly concern itself with its subject matter. Which brings us to category B.

Category B work is viewed unenthusiastically by Lighthill, and defended with vigour by others. One line of defence is to call attention to developments in programming technology which it has stimulated, and to other insights to which it has led. In any venture into the history of ideas one is on dangerous ground, but in considering this kind of argument the risk has to be taken. I do not believe that the case can be made by considering programming technology. Structured programming has no dependence on AI, and the handling of complex low-level operations in terms of smaller numbers of higher-level notions has been taken to its highest development by people whose view of AI is no more favourable than Lighthill's. Backtracking is a programming technique of much antiquity. The embodiment of knowledge in procedures is a year or two younger than the act of programming; its descriptions for the plain man is that, when looking up a table you sometimes find the address of a program to compute the value you want rather than directly being given the value itself. List-processing is a technique for burying store-management problems, excellent for rich people with complicated programs to write.

It is beyond contention that AI research has led to a great deal of excellent work in packaging these techniques attractively and embodying them in programming languages (some of which, for example, LISP, are of much interest as languages); it is a standard progression for frequently used facilities to start as library routines and end as language features.

However, the ideas did not all originate in AI, any more than did the content (though perhaps not the phraseology) of the maxim 'heterarchy not hierarchy'. The general inapplicability of strict hierarchical models, despite their seductive clarity has not merely been known but explicitly recognised as an important point by many people for a long time. To be explicit about a few aspects in my own experience: Library Automation – middle 1950's; Taxonomy – early 1960's; computer filing systems – middle 1960's. To the small extent that so vague and general a maxim can be said to be a discovery, it is one to which AI has contributed little.

Professor Sutherland remarks, in the course of a defence of category B work along these lines, that 'one recent insight derived from Basic research on AI is that in interpreting the meaning of any complex input, it is impossible to use a rigid step-by-step procedure'. Leaving aside the perhaps captious comment that this means that computers cannot do the job at all, it is emphatically **not** a recent insight that you cannot finish with the syntax before starting the semantics. The present writer first encountered it, in a computational linguistic context, sixteen years ago when it was not new.

This comment on a point to do with language processing leads naturally to others. Lighthill cites the understanding of natural text as one of the prime examples of the combinational explosion, and so it is. He, and also Sutherland, do in my opinion underestimate the contribu-

tions which have come from activities which are (or were not) called AI. Most people writing on such subjects tend to dismiss Machine Translation not only as a technological failure, which it was, but as an intellectually totally negligible activity, which it was not. The emphasis on exact algorithms rather than vague descriptions, the above-mentioned importance of mixing syntax with semantics, the use of heuristic devices to shorten searches in great collections of semantic data, were all studied and recognised as important. Linguistics proper, without technological objectives, has made vast progress in schematisation, exactness of description, and theoretical understanding. General interest in linguistics and in particular in rule-based (alogrithmic) approaches to it, have contributed much to the intellectual climate in which AI work is done, though not always in a directly recognised manner. It is a great mistake to extract those products of other enquiries which have been found helpful by AI workers and suppose that they are AI inventions.

In sum, I do not believe that one can justify category B work by its side effects. What of its main thrust? Whether there is any middle ground between studying intelligent behaviour as an attribute of people or animals on the one hand, and making machines do complicated and useful things which used to need people, on the other, it can easily lead to a sterile philosophical debate unless we say 'yes', on the grounds that we can seek to make machines do complicated and useless things.

It is perhaps best left as a personal position that it is unlikely that work in AI which is neither motivated by a stateable technological aim nor by definite study of how people work will be productive. In short, I agree with Lighthill. It will always be a matter of judgement whether a particular proposed activity is a real technological one or a badly over simplified one, or whether such an activity, if technological in aim, is naive in method or interesting in method, or whether a certain technological aim has given rise to a valid goal for general enquiry. Perhaps what is most lacking in much category B work is the chain of reasoning starting from a technological aim and finishing with a need to solve a specific general problem.

'Artificial Intelligence' is a rather pernicious label to attach to a very mixed bunch of activities, and one could argue that the sooner we forget it the better. It would be disastrous to conclude that AI was a Bad Thing and should not be supported, and it would be disastrous to conclude that it was a Good Thing and should have privileged access to the money tap. The former would tend to penalise well-based efforts to make computers do complicated things which had not been programmed before, and the latter would be a great waste of resources. 'AI' does not refer to anything definite enough to have a coherent policy about in this way.

A final caution: like the majority of contributors to this paper symposium on AI, I am not an expert in any of the activities which come under its rather ill-defined umbrella. Amongst the many features, good and bad, which AI shares with Machine Translation is the fact that non-practitioners have strong views about it.

2 Comment by Professor H. C. Longuet-Higgins, FRS.

Theoretical Psychology Unit, University of Edinburgh.

To my mind Sir James Lighthill's most valuable contribution to the current debate on artificial intelligence has been to raise searching questions about the proper justification of the subject. We should, he suggests, ask about any piece of work whether its primary objectives are technological or scientific. If technological, such as the automatic exploration of the planets or the mechanical translation of Chinese into English, are such aims realistic in relation to our present knowledge and justifiable in economic terms? If scientific, then what science or sciences are likely to be enriched?

The need to ask such questions becomes only too apparent when one studies certain recent pronouncements on the subject (or subjects?) of artificial intelligence/machine intelligence. In the Computing Science Review recently published by the SRC the aims of 'machine intelligence' are seen as bluntly technological, though the pious hope of achieving them 'by the formulation of general principles' makes a hasty genuflection to scientific respectability. There is no further mention of any such principles either in the main body of the review or in the appended report of the Long Range Research Panel; all we find is an unlikely assortment of subjects grouped together under the heading 'machine intelligence' for no better reason than 'none of [them] seem to the Panel to demand the study of human thought or perception'.

The subjects in question are: computational logic, real scene analysis, picture processing, 'the use of the robot as an analytical tool', and 'the acquisition of organised information in computers and interpretation of descriptive material'. It is highly dubious whether either computational logic or real scene analysis is likely to get anywhere without due attention to our processes of thought and perception; but in any case, if such a negative criterion is adopted for what counts as 'machine intelligence', it is difficult to see why that subject should exclude analytical geometry or analytical chemistry, which have at least as good a claim to be regarded as 'analytical tools'.

The poverty of such arguments for regarding machine intelligence as a priority area in computing science must have become plain to Sir James as soon as he undertook his penetrating survey of the field. Insofar as machine intelligence projects are basically technological, should they not be judged by the same criteria as one would apply to any piece of development work in advanced automation? Of any such project one should ask: first, what exactly is it intended to achieve, secondly, what material resources would it demand, and thirdly, what are its chances of success? Lighthill's shrewd and comprehensive critique of the technological achievements and ambitions of artificial intelligence needs no recapitulation, nor does his scepticism about the defensibility of robotics as a technological enterprise. It is only when one looks at the scientific case for artificial intelligence studies that differences of opinion seem to arise.

Sir James places in his category 'C' all the artificial intelligence work which he regards as scientifically promising, and refers to this category as 'Computer based studies of the central nervous system'. In so doing he aligns himself with those of us who hold that the main justification for artificial intelligence is the light it can throw upon human intellectual activity. But his chosen heading, and some of his later remarks, indicate that he attaches more significance to work on the 'hardware' of the brain than to work on its 'software'.

This is the only point on which I want to take issue with him. He is, of course, perfectly right in saying that anyone who is developing network models of the brain had better work within the constraints imposed by our knowledge of its anatomy and physiology; it would be foolish for an engineer to speculate about the circuitry of a computer when he could perfectly well open it up and look inside.

But the hardware of computers is very far from being the only matter relevant to their functioning. In order to understand how a computing system works one must enquire into the logic of the system software and the semantics of the programming languages in which the system can be addressed. The corresponding questions about human beings are those asked by the science of psychology – though admittedly, psychological theories seldom attain a degree of sophistication worthy of their subject matter. An outstanding exception to this stricture is the science of linguistics, and perhaps it is no coincidence that the most impressive achievement of artificial intelligence to date is a working model of the comprehension of natural language.

I would go further and hazard the prediction that for some time to come the most valuable work in artificial intelligence will be that which attempts to express, in the form of computer programs, abstract theories of our various cognitive faculties, rather than mathematical models of the brain itself – this in spite of some excellent recent work on the possible role of the neocortex as a classifying device. This view is based not only on the obvious vitality of current artificial intelligence work on language and vision, but also on an evident dissatisfaction among psychologists with the naive stimulus-response theory of behaviour as it has been applied to human beings.

It is now plain that a central problem in cognitive psychology is to understand how our knowledge is represented and deployed, and the computer program is the only medium which at present offers us the possibility of formulating adequately sophisticated theories of cognition. The elimination of inadequate theories is no longer the main problem; the defects of a programmed theory become immediately apparent as soon as it is run on a computer.

In short whatever the technological prospects of artificial intelligence, its principal scientific value, in my view, is that **it sets new standards of precision and detail in the formulation of models of cognitive processes, these models being open to direct and immediate test.**

The question 'What science or sciences are likely to be enriched by artificial intelligence studies?' can now receive a provisional answer,

namely 'All those sciences which are directly relevant to human thought and perception'. These 'cognitive sciences' may be roughly grouped under four main headings:

1 **Mathematical** – including formal logic, the theory of programs and programming languages, the mathematical theory of classification and of complex data structures.

2 **Linguistic** – including semantics, syntax, phonology and phonetics.

3 **Psychological** – including the psychology of vision, hearing and touch, and

4 **Physiological** – including sensory physiology and the detailed study of the various organs of the brain.

Perhaps 'cognitive science' in the singular would be preferable to the plural form, in view of the ultimate impossibility of viewing any of these subjects in isolation. Indeed artificial intelligence studies are beginning to offer interesting suggestions as to how our various modes of experience might be logically related.

Finally, perhaps one should say a word about the main point of disagreement between Lighthill and Sutherland. Professor Sutherland's redefinition and reinstatement – of Lighthill's category 'B' as 'basic artificial intelligence' has my sympathy, because although I hold no particular brief for 'bridging activities' as such, I do think that there is a place in artificial intelligence for studies which are addressed to the general problems which have been found to recur in many different areas of cognitive science. The mathematician's ability to discover a theorem, the formulation of a strategy in master chess, the interpretation of a visual field as a landscape with three cows and a cottage, the feat of hearing what someone says at a cocktail party and the triumph of reading one's aunt's handwriting, all seem to involve the same general skill, namely the ability to integrate in a flash a wide range of knowledge and experience. Perhaps Advanced Automation will indeed go its own sweet way, regardless of Cognitive Science; but if it does so, I fear that the resulting 'spin-off' is more than likely to inflict multiple injuries on human society.

3 Comment by Professor D. Michie

Department of Machine Intelligence and Perception, University of Edinburgh.

> *'Two contrary attitudes are common. In the first place there is a widespread, although mostly unconscious, desire to believe that a machine can be something more than a machine, and it is to this unconscious urge that the newspaper articles and headlines about mechanical brains appeal.*
>
> *'On the other hand, many people passionately deny that machines can ever think. They often hold this view so strongly that they are led to attack designers of high-speed automatic computing machines, quite unjustly, for making claims, which they do not in fact make, that their machines have human attributes'.*
>
> M. V. Wilkes 'Can a machine think?' in *Discovery* May, 1953.

Sir James Lighthill's report speaks of 'the ABC of the subject', categorising it as follows:

A – Advanced Automation

B – Building Robots

C – Computer-based CNS research

The report regards 'A' and 'C' as worthy activities which, however, have made disappointing progress. 'B' is regarded as unworthy, and as having made very disappointing progress indeed. 'B', it should be noted, is really used in the report to denote any experimental programming which lacks obvious opplication to either 'A' or 'C'. Thus computer chess is included in 'B' whereas robot parcel-packing is put into 'A'.

Most people in AI who have read the report have had the feeling that the above classification in misleading. Sir James has arrived at his position by interpreting AI as consisting merely of outgrowths from a number of established areas, viz.:

A as an outgrowth from control theory,

B as an outgrowth from science fiction,

C as an outgrowth from neurobiology,

These interpretations are remote from those current in the field itself. A number of questions accordingly pose themselves, including the following:–

1 Was this report based on as thorough a survey as it should have been? In particular, was opportunity taken to invite the views of the international leaders of the field?

2 How successful has the author been in overcoming the difficulties inherent in his inexperience of the field, and in putting aside his own professional biases?

3 Has accepted practise been followed in documenting subjective opinions wherever possible and in providing factual sources and references which others can check?

4 What is the validity of the 'A B C' classification? Would the computing science community accept it?

5 Are the report's assessments of work in the 'B' category — 'Building Robots' — intended to apply to experimental robotics conducted in the United Kingdom? If so, should not the author

i have said so plainly,

ii have asked to see the experimental robotics work during his visit to Edinburgh?

International opinion not consulted

The first of these questions is so critical as to merit a brief note to the effect that the leading American workers, such as McCarthy, Minsky, Nilsson, Raphael and Robinson, were not in fact consulted. The appearance of their names in the list of fifty given in the report's third paragraph derives from the fact that the author has read scientific writings of theirs not that he invited their opinions. Since the field in question was pioneered in the United States of America, which supports to this day an effort on at least twenty times the scale of that in the United Kingdom, it is well to bear in mind this fact when assessing Sir James' evaluations.

Intelligence Theory

Space does not allow a review of the remaining questions of the above list; a detailed critique is available elsewhere.* Instead I will briefly indicate two themes which arise from Lighthill's implicit question: 'If you throw A and C away, what is left, if anything?' Lighthill's answer is 'Building Robots'. An alternative answer, which many will prefer, is 'Intelligence Theory'. By this we mean attempts to systematise the design principles of intelligent systems wherever they may be found, whether in the 'A' or 'C' application areas.

Having fixed on 'Building robots', Lighthill paints a picture of this pursuit which must strike those actually engaged in experimental robotics as somewhat unfamiliar. In studies of actual robot work the role of the equipment is plainly seen as test gear for putting certain types of theoretical ideas to experimental test. A pertinent parallel is the building of wind-tunnels as an aid to aero-engineering — as illustrated in the *figure* on page 40.

This *figure* brings into relief the reason why Building Robots is an unhelpful choice for the role of Bridge between 'A' and 'C'. It is surely more fruitful, if one seeks inter-disciplinary connections, to choose a common body of theory rather than to seize on a piece of laboratory equipment. One feels that Sir James would be among the first to agree that to speak of Building Wind-tunnels as the Bridge between aero-engineering and the study of bird flight, would direct attention away

* D. Michie (1972) 'On first looking into Lighthill's 'Artificial Intelligence' report' (unpublished).

Broad subdivisions of two fields of enquiry according to 'theoretical-experimental' and 'technological-biological' classifications.

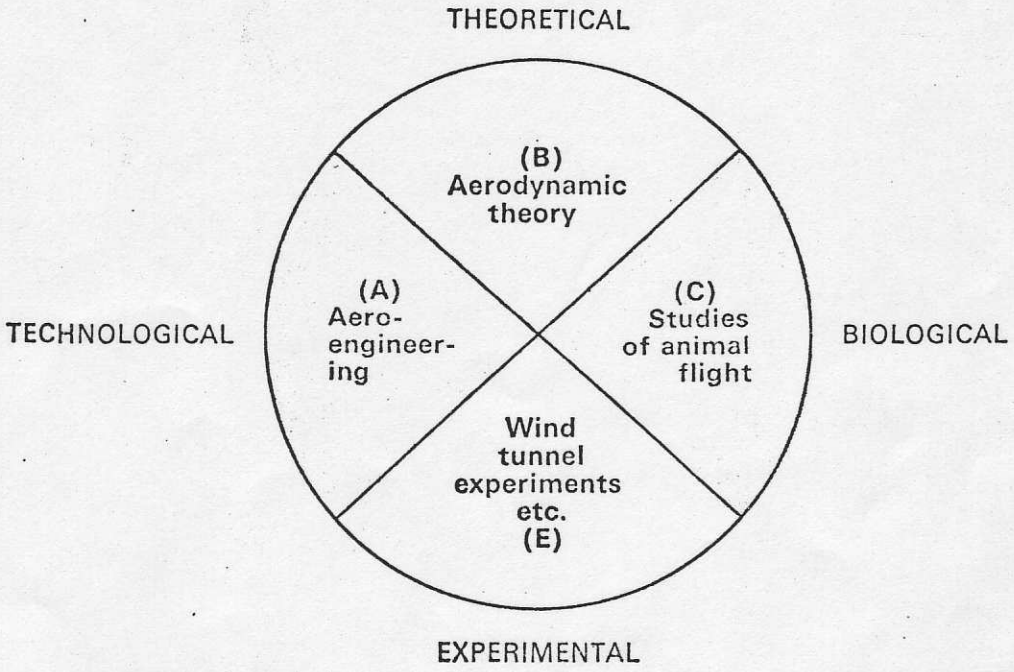

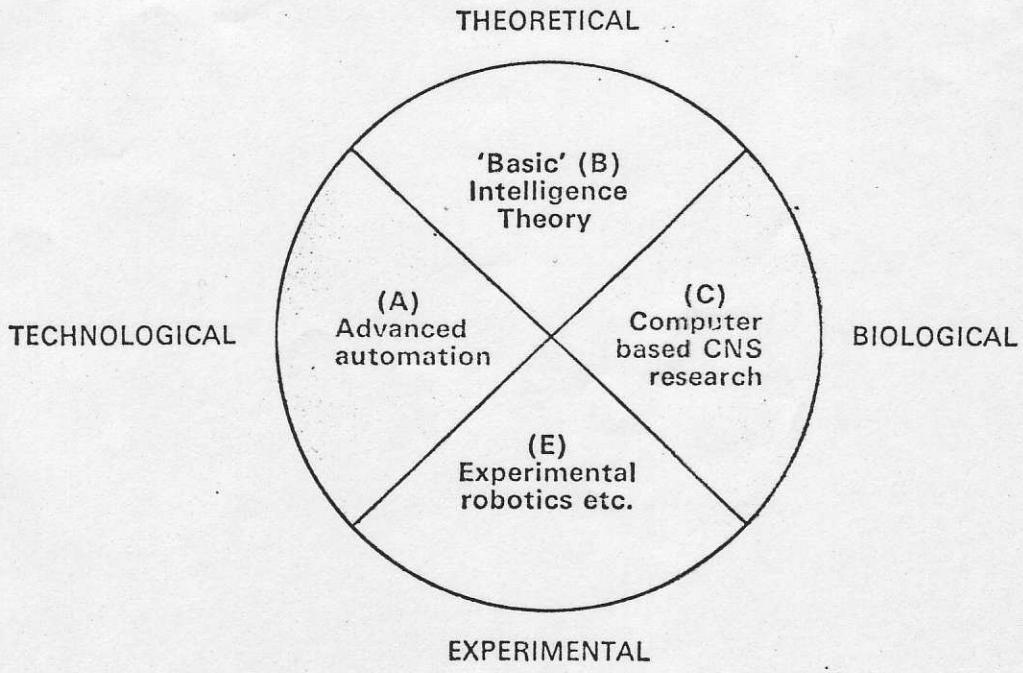

from the true bridge, namely the science of aerodynamics. The equivalent science in the case of AI is at a primitive stage. It is the hope of every AI professional to contribute in some way to bringing the required theory into being. This, as I see it, is the burden of Sutherland's re-definition, in his contribution to this symposium, of 'B' as standing for 'Basic'.

Encouragement of research in machine intelligence

On this note I would like to leave Sir James Lighthill's interesting and imaginative review and to mention an assessment of a more home-spun quality: the report of SRC Computing Science Committee's long-range panel, published in 'Computing Science Review'. This panel, composed of computer professionals, considered the machine-oriented part of Artificial Intelligence (ie the 'A+B' part) and recommended that special encouragement should be given to this field. However, it is evident to those who work in the field that it would be helpful if a clear and concise statement were given of its goals and methodology. The style of Sir James Lighthill's report suggests that there is a lack of understanding in some quarters, and without this there is a reluctance to recommend significant expenditure. The status and position of the subject are particularly clear at the moment and it is, therefore, opportune that a statement should be made to avoid any further misunderstandings.

The subject, in so far as it comes within the Computing Science Committee's realm of interest, is concerned with machines, and in particular computers, displaying characteristics which would be identified in a human being as intelligent behaviour. Perhaps the characteristics which are most important are those of learning and problem solving. The applied benefits which may be gained from work in this field could bring considerable economic benefit to the country. They are two-fold:

a To relieve the burden at present on the systems analyst and programmer in implementing applications:

b To enable new and more complex applications to be undertaken in this country in competition with work elsewhere.

These are the long-term advantages and to this end work is proceeding on a number of detailed problems, including the following:

I Automatic assembly and other robotic applications
Mass spectogram analysis
Chemical synthesis planning
Assembly-line balancing

II Language-understanding systems
Semi-automatic programming (*ie* 'teachable systems') and ultra-high level programming languages (like PLANNER, SAIL, CONNIVER)

Group I represents useful applications. Group II represents the subject's own special contribution, independent of specific applications to computer science. This lies in making it more possible for the user to get computing systems to 'understand' what he means.

Many good scientists have been involved in this field and their work has resulted in the development of techniques and methods of wider use, for example:

List-processing was originally devised by Newell, Shaw and Simon for AI work and first implemented in their IPL language.

The incorporation of conditional expressions into ALGOL 60 was McCarthy's suggestion derived from his work on LISP, itself inspired by the needs of Artificial Intelligence work.

The POP-2 language, now implemented on 5 main hardware ranges was specifically developed for AI work, but subsequently shown to be of wider utility.

In fifteen years of struggle towards 'language understanding', striking advances have been scored (Bobrow, Winograd, Woods).

Some of the search and associative techniques used by programmers and operations research workers have been initiated in AI, and assimilated without awareness of their origin.

The problems that have been mentioned above are practical problems. Abstracting from these, and observing the methods of solution, workers in the field have been able to define general principles for intelligent systems. This work has made some progress and the following theorems and methods have been developed.*

Problem-solving

Theorems of minimality and completeness of various algorithms for heuristically guided search.

Methods of pruning search trees in special situations: Plausibility analysis; alpha-beta pruning.

Recursive formation of sub-problems as in Newell and Simon's 'General Problem Solver'.

Application of theorem-proving ideas in problem-solving.

Studies of problem-representation.

Recognition

Various methods of feature extraction and interpretation for visual data.

Use of semantics to disambiguate linguistic analysis.

Matching of descriptions represented as directed graphs (*eg* 'hierarchical synthesis').

* Some of these are reviewed in a *Nature* article 'Machines and the theory of intelligence' 23 Feb. 1973.

Learning

Adaptive learning *via* parameter-optimisation.

Rote-learning techniques.

Formation of new concepts from examples and counter-examples.

Inductive generalisation.

Even so incomplete a list as the above puts into perspective the importance of examining particular problems in depth (such as chess-playing or those involving robots) so as to investigate how to bring the above functions to bear in an integrated fashion. They are but experiments which may be used to derive or test theories. At this early stage of innovation the overwhelming benefit to be derived from a given experimental study lies in its role as a forcing function for new programming techniques and tools. The field is so difficult and the choice of the right problem at the right moment so much part of the art of enquiry that this should be left to the research workers themselves. They should be judged by their success or otherwise in advancing the state of computer programming, and in introducing and testing computer languages of greater expressive power.

Footnote on Sutherland's commentary

Sutherland's otherwise admirable analysis contains two expressions of view with which exception must be taken, namely 1, that AI should not be handled by the Engineering Board of SRC and 2, that AI research in Britain is in a bad way.

1. A reasonable approach would surely be to distinguish A-oriented and C-oriented poles of the subject and to provide for the first under the Engineering Board and for the second under the Science Board. Since important contributions continue to be made by computer scientists ignorant of psychology and brain science, and by psychologists ignorant of computer science, it would avoid embarrassment, and reflect scientific reality, to make separate provision.

2. Sutherland's proposition that AI research in Britain is in a bad way deserves to be vigorously challenged. But it is inappropriate for me, as founder of the longest-established British research group, to be the one to do this. A better corrective can be obtained from assessments by authoritative outside observers, such as that by Dr. Nils Nilsson* of the Stanford Research Institute's Artificial Intelligence Centre and author of the graduate textbook 'The problem-solving methods in Artificial Intelligence'.

What is to be done

We are in the embarrassing situation in Britain that in order to carry out significant work over the next few years in the context of international competition, it will be essential to import American machines — speci-

* 'An outsider's view of the Experimental Programming Unit at Edinburgh University', obtainable from School of Artificial Intelligence, Edinburgh.

fically the DEC System 10.* I think that everybody would be happier about the case for allowing an American importation now if steps were at the same time taken **to see that British AI research never found itself in such a predicament again.** What would have to be done if this desirable state of affairs were to be brought about?

Why does the need arise? It is not only, or primarily, because of the superiority of the architecture of the DEC System 10 for AI-type uses. The over-riding consideration is access to the rapidly accumulating fund of AI-oriented software and applications programs in the big American laboratories. The key to the situation is the absence of software for British machines, either present or 'new range', suitable for AI work, which has its own very peculiar needs. These needs are peculiar. It is hardly more sensible to speak of 'making do' with, say, general scientific software developed without reference to AI than to suggest that, say, plasma physicists short of experimental fusion equipment should 'make do' by borrowing linear accelerators from the particle physicists!

The kind of software development needed if AI workers are ultimately to be put in business as users of the new range of British computers (I do not necessarily intend this phrase to be exclusively confined to ICL) comes under two headings:

1. Development of experimental operating systems, compilers and packages, as has been done in a small way on the ICL 4130 at Edinburgh. But the new effort should aim to embrace the entire 'standard range' of facilities which every AI worker should be entitled to take for granted – LISP, POP-2, SNOBOL, QA4, PLANNER, CONNIVER, etc. etc., and, ultimately far more important ,'leap-frogging' into the future both by adapting the latest advances of AI research work where appropriate, and by innovation within the R & D effort itself. Also to be considered are operating system features for handling 'funny peripherals' (experimental robots, speech input devices, *etc.*) and basic packages for 'front end' functions such as, say, video and speech input, robot control functions, language pre-processing. In addition the design and development of advanced peripherals (eg for robotics) should be regarded as an integral part of the job although (as with software) the more standard aspects of instrumentation should be contracted out to industry wherever possible.

2. 'Communality aids' whereby new research programs and software developed in overseas laboratories can be made immediately available on demand for British research workers to test out and either accept or reject as tools for their own needs. Communality can be achieved by various means and these means will vary according to the nature of the case, but they include software/hardware interfaces to the ARPA net, and emulation (for example by microprogramming) of the 'donor' machine from which the program is to be adapted.

If a fully-fledged development project is to be got up to full speed by around 1977 then forward studies could usefully be started now. **It is already obvious that early installations of a PDP-10 in an active centre of British AI research is a precondition if these studies are**

* formerly known as the PDP-10

to develop fruitfully, since immediate access to the latest AI research materials (and intimate contact with advanced AI research) will be as essential to the specification and development of new research facilities as it is for those who will later be using them. Until the new facilities exist, the only point of access and contact will be through British groups equipped compatibly with their American counterparts, — *ie*, with PDP-10s.

In this field successive workers in a given area should be able to stand on the shoulders of their predecessors through the medium of successive contributions to a common stock of new language aids and library packages. This will not happen unless someone makes it his business continually to scoop in what is new and useful and build it into a properly documented and integrated system. The level at which the British AI community will be able to contribute in the late 1970s, as judged by competitive international standards, will be crucially affected by the sophistication of the available software.

Is this a rebuttal of Lighthill or a request for expansion?